

NBER WORKING PAPER SERIES

FROM FLAT TO FAIR? THE EFFECTS OF A PROGRESSIVE TAX REFORM

Nicolas Ajzenman
Guillermo Cruces
Ricardo Perez-Truglia
Darío Tortarolo
Gonzalo Vazquez-Bare

Working Paper 33286
<http://www.nber.org/papers/w33286>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
December 2024

No additional disclosures needed. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research, nor the International Bank for Reconstruction and Development/World Bank and its affiliated organizations, or those of the Executive Directors of the World Bank or the governments they represent.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2024 by Nicolas Ajzenman, Guillermo Cruces, Ricardo Perez-Truglia, Darío Tortarolo, and Gonzalo Vazquez-Bare. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

From Flat to Fair? The Effects of a Progressive Tax Reform

Nicolas Ajzenman, Guillermo Cruces, Ricardo Perez-Truglia, Darío Tortarolo, and Gonzalo Vazquez-Bare

NBER Working Paper No. 33286

December 2024

JEL No. C93, D31, H24, H26, H71

ABSTRACT

This paper investigates the impact of a progressive tax reform on tax compliance. We leverage a major progressive tax reform in a large Argentine municipality. First, we use a quasi-experimental design to estimate the causal effect of changes in a household's own tax rates on its tax compliance. Second, we utilize a large-scale natural field experiment to examine whether, holding a household's own tax rates constant, tax compliance is influenced by the tax rates of poorer or richer households. We find that reducing taxes for poorer households increases their compliance, while increasing taxes for richer households decreases their compliance. When poor households learn about the tax hike on the rich, this increases their perceived fairness of the tax system and their tax compliance. When rich households learn about the tax cuts for the poor, their perceived fairness increases significantly, but their compliance, if anything, goes down. Leveraging another reform (and another field experiment) that took place a year later, we show that both the quasi-experimental and experimental findings replicate. Our evidence highlights that tax compliance depends not only on a household's own tax rate but also on its perception of the broader tax schedule. Our findings also highlight the gap between stated and revealed preferences for redistribution. Lastly, we conduct a counterfactual analysis to illustrate the implications of our findings for the design of tax policies.

Nicolas Ajzenman
McGill University
855 Sherbrooke St W
Montreal, Queb
Canada
nicolas.ajzenman@mcgill.ca

Darío Tortarolo
The World Bank
1818 H Street, NW
Office MC3-361
Washington, DC 20433
dtortarolo@worldbank.org

Guillermo Cruces
CEDLAS
Univesidad Nacional de La Plata
Calle 6 entre 47 y 48
La Plata, Argentina
gcruces@cedlas.org

Gonzalo Vazquez-Bare
Department of Economics
University of California, Santa Barbara
Santa Barbara, CA 93106
gvazquez@econ.ucsb.edu

Ricardo Perez-Truglia
University of California, Los Angeles
110 Westwood Plaza
Los Angeles, CA 90095
and NBER
ricardotruglia@gmail.com

A data appendix is available at <http://www.nber.org/data-appendix/w33286>

A randomized controlled trials registry entry is available at AEARCTR-0010738

1 Introduction

Progressive tax schedules are widespread throughout the world and play a crucial role in the redistribution of income (Piketty and Saez, 2007; Saez and Zucman, 2019a).¹ There are, however, significant differences in the degree of tax progressivity between countries and even across different taxes within a given country (Fisher-Post and Gethin, 2023). A natural hypothesis for the adoption of progressive tax schedules is that individuals care not only about how policies will affect them, but also about how they will impact others, for example, through other-regarding preferences or social comparisons (e.g., see Stantcheva, 2021). For instance, wealthy households may tolerate a reform that increases their tax burden if it advances fairness objectives, such as supporting disadvantaged households. Although it is such a pillar of modern tax systems, real-world evidence on the factors behind the support for tax progressivity is scarce. This study addresses this gap by analyzing how taxpayers respond in practice to a real-world progressive tax reform.

We combine experimental and quasi-experimental methods to show that increasing tax progressivity has substantial effects on tax compliance, and that those effects vary depending on how the reform affects the individual taxpayer and their peers. Our findings offer valuable insights into individuals' true attitudes towards progressive taxation. Individuals may state that they prefer more progressive taxes (Tarroux, 2019), but talk is cheap – do they put their money where their mouths are? The behavioral responses to a progressive reform give us insights on the attitudinal and practical consequences of such reforms and on the determinants of the support for progressive taxation in general. Our results have important implications for designing revenue-neutral tax reforms and for their support, particularly in contexts where tax enforcement is limited.

We hypothesize that progressive taxation could affect tax compliance through two distinct channels. The *own-rate* effect posits that a taxpayer's tax compliance may respond to a change in their own tax rate, irrespective of what happens to the tax rates of other households. For instance, lowering the taxes on a poor household may increase its tax compliance, while increasing the taxes on a rich household may lower its tax compliance. The *cross-rate* effect refers to the fact that, holding constant its own tax rate, a taxpayer's compliance may also depend on what the household perceives about the tax rates of other households. In the context of a progressive reform, a poor household may change its tax compliance after finding out that, in addition to lowering its own tax rate, the government is increasing taxes on the rich. Likewise, a rich household may change its tax compliance upon learning that the government reduced tax rates for the poor, over and above the effect of an increase in its own tax rate. Ex-ante, cross-rate effects could

¹For example, according to Fisher-Post and Gethin (2023), progressive tax systems reduce inequality by 5% to 15% in Western Europe and the United States – for more evidence, see also Lustig (2023).

be positive or negative, depending on the prevailing social preferences and how different groups of taxpayers are affected. For example, some taxpayers may react negatively to increased progressivity if they believe that it would be fairest for everyone to pay the same rate. On the other hand, if taxpayers want more progressive taxes, finding out about a progressive reform may boost their tax morale.

We study a progressive tax reform on property taxes implemented by Tres de Febrero, a major municipality in the province of Buenos Aires, Argentina (and a suburb of the nation's capital). The tax is levied monthly on the assessed values of properties, with an average tax rate of about 3%. Tax revenues are used to fund the provision of basic services such as street lighting and urban sanitation. With guidance from our research team, the government designed and implemented a progressive tax reform in January 2023. Properties at the bottom 40% of the value distribution (hereafter referred to as *poor households*) received a tax cut, while those in the top 22% (*rich households*) faced a tax increase.² Tax rates for properties in the middle part of the distribution (*middle-income households*) remained unchanged. The magnitude of the reform was significant: on average, taxes decreased by about 27% for the poor while they increased by about 11% for the rich.

The setting in which this reform occurred has two advantages in addressing the research question at hand. First, there is substantial wealth inequality, and thus there is ample scope for redistribution through taxation. For instance, the Gini coefficient for the greater Buenos Aires (which includes Tres de Febrero) was 0.404 (CEDLAS, 2024), which is comparable to that in cities such as Miami, Lisbon, and Brussels (OECD, 2018). This inequality is reflected in home values too. For example, in Tres de Febrero, a property in the 90th percentile of assessed value is worth 7.8 times as much as a property in the 10th percentile. The second advantage is that, as in most low- and middle-income countries, there is imperfect tax enforcement and thus substantial noncompliance. Thus, there is scope for households to increase or decrease their tax compliance in response to the reform. For example, in the year prior to the reform, the average probability of making timely payments (i.e., within three months of the due date) was 47%.

We identify the *own-rate* effects of the reform using a two-cutoff Regression Discontinuity (RD) design that leverages two sources of identification. First, the reform involved two sharp discontinuities: households with property values below a predetermined threshold (poor households) were hit with a tax cut, while households with property values above a different threshold (rich households) were hit with a tax hike. Properties in the middle group did not experience any tax change. Moreover, we leverage the timing of the reform

²There was a further tax increase for properties in the top 6% of valuations, but we lack statistical power to analyze the impact of this additional change separately. Therefore, the top 6% of the valuations are excluded from the whole analysis. For this reason, what we refer to as the top 22% richest properties technically correspond to the percentiles 73rd through 94th.

as a second source of causal identification: we compare the evolution of the outcomes right before versus right after the reform took effect.

The RD estimates reveal significant own-rate effects of the reform on tax compliance, which we can summarize as behavioral elasticities. A 1% reduction in the tax rate for the poor increases their compliance by 0.17%, implying an elasticity of $\varepsilon_{poor}^{own-rate}=0.17$. Conversely, a 1% increase in the tax rate for the rich reduces their compliance by 0.36%, corresponding to an elasticity of $\varepsilon_{rich}^{own-rate}=0.36$.³ While we do not have data to disentangle the two, our preferred interpretation is that the own-rate effects are likely a mix of “price” effects and tax morale effects. For example, maybe some poor households wanted to comply all along but were not able to afford their tax payments. The “price” effect means that by lowering their tax obligations, their taxes become more affordable and thus they become more likely to pay them (e.g., see [Bergeron et al., 2024](#); [Brockmeyer et al., 2023](#)). Likewise, by increasing taxes on the rich, they become less affordable, and thus rich households become less likely to pay. However, the effects of the own rate could be due, at least in part, to tax morale. For instance, in response to the tax cut, the poor may want to reciprocate the government’s gesture by paying their taxes, or they may increase their tax compliance because their perceived fairness increases. Likewise, a tax hike on the rich may push them to reciprocate by stopping paying their taxes, or it may lower their perceived fairness of the system.

These RD estimates are robust to a host of sensitivity checks. Most notably, we provide two key falsification tests: we find null effects when using placebo thresholds and when using a fake date for the policy change. Moreover, to see if the results replicate, we leveraged a second progressive tax reform that took effect a year after the first reform, in January 2024. Since the reform also used thresholds to determine which households would experience a tax cut and which ones would experience a tax hike, we can reproduce the RD analysis exactly as for the first reform. We find estimates that are both qualitatively consistent and in the same order of magnitude.

To identify the *cross-rate* effects of the reform, we conducted a pre-registered, large-scale, natural field experiment, in which we randomized information about the progressive nature of the tax reform. As part of regular communication with taxpayers, the municipality sent letters to almost the entire universe of taxpayers (about 100,000 households). We embedded an information provision experiment in the January 2023 mailers. We randomly assigned each subject to receive one of two types of letters. The first type included information solely on how the reform changed the household’s own tax rate. The second type included additional information about how the reform affected the tax rates of other households. For instance, a poor household could receive either a (control) letter

³Since their tax rate was not affected by the reform, we cannot estimate an *own-rate* effect for middle-income households.

detailing the tax cut for poor households or a (treatment) letter that also highlighted the broader reform, that is, the tax hike for rich households. This setup ensured that information about changes to a household’s own tax rate was held constant while experimentally varying awareness about how the reform affected other households.

Using administrative records, we can estimate the impact of treatment on subsequent tax compliance. In addition, to provide complementary evidence on the causal mechanisms at play, we collected survey data. More precisely, we invited a small subsample of taxpayers with email addresses on file to respond to a short online survey. Approximately 2,150 households responded to the survey, for a response rate of 16.4%. One potential concern with information-provision experiments is that subjects may not pay attention to the information provided to them or that the control group may find out about the information through other means. To address these concerns, the survey included a question designed to assess awareness of the progressive tax reform: subjects received a list of recent policies and were asked to indicate which ones they had heard of. With this outcome, we can measure whether the treatment actually had a significant effect on awareness of the progressive reform. Finally, to investigate the tax morale mechanism, one survey question elicited perceived fairness of the tax system, using a subjective scale from 0 (very unfair) to 10 (very fair).

We first discuss the cross-rate effects for poor households. Using survey data, we show that the treatment had a significant effect on awareness of the reform. In the control group, a minority of households are aware of the reform, while the awareness is about 50% higher in treatment households. Using administrative records, we show that the treatment had a positive and significant effect on tax compliance. That is, poor households’ tax compliance increases when they learn about the progressive nature of the reform. Compared to poor households who received the control letter, those treated with information about the progressivity of the reform were 0.95 percentage points (p-value = 0.005) more likely to pay their due taxes during the following three months. We estimate a cross-rate elasticity of $\varepsilon_{poor}^{cross-rate}=0.11$ for poor households: a 1% increase in the (perceived) taxes of the rich increases their own tax compliance by 0.11%. To put this magnitude in context, we can compare it to the own-rate effect: a 1% increase in the tax rate of the rich has the same effect on the compliance of a poor household as lowering the poor household’s own tax rate by 0.66%. While the magnitude of this cross-rate effect is already substantial, this is just an intention-to-treat estimate, so the treatment effects on the treated could be twice as high.⁴ By construction, since the household’s own taxes are held constant, the cross-rate effects cannot be attributed to “price” effects. Instead, our preferred interpretation is that the cross-price effects are driven by the tax morale

⁴For example, a significant fraction of the households may not have received the letters, or they have not read them carefully.

mechanism. Indeed, we use the survey data to provide some direct evidence in support of this mechanism: when informed about the tax hike on the rich, poor households became significantly more likely to rate the tax system as fair.

Next, we discuss the cross-rate effects for rich households. As with poor households, we find that the treatment increased rich households' awareness of the progressive tax reform. Moreover, when the rich find out that the poor are getting a tax cut, they are more likely to believe that the tax system is fair. Although their perceptions of the tax system's fairness increased by about the same degree as for the poor, we do not see a positive effect on tax compliance for the rich. If anything, the rich become *less* likely to pay their taxes after finding out that the poor experienced a tax cut, although the effect is close to zero in magnitude ($\varepsilon_{rich}^{cross-rate}=0.01$) and statistically insignificant. We estimate the cross-rate effects for middle-income households. Although they do not see a change in their own tax rates, finding out about the tax hike on the rich and the tax cut on the poor may still affect the compliance of middle-income households. The results for the middle-income households mimic the results for the rich households: the treatment had a significant positive effect on the awareness of the progressive tax reform and on the perceived fairness of the tax system; however, the effect on tax compliance is close to zero and statistically insignificant.

The results of the field experiment are robust to a host of sensitivity checks. Most importantly, we conducted an event-study analysis showing that the timing of the effects of the treatment coincides exactly with the timing of the mailing intervention. Moreover, to see if the results replicate, we conducted another pre-registered field experiment in January 2024, leveraging the second progressive tax reform that took place at that time. The results for the 2024 experiment are largely similar to the results from the 2023 experiment. Lastly, to assess how surprising these results were, we conducted a prediction survey with 39 academics with relevant research experience. After receiving a description of the experiment, these experts were asked to forecast the treatment effects. The survey revealed a lack of expert consensus, with significant variation in predictions and most experts expressing uncertainty about their forecasts. Only a minority of experts predicted effects close to our experimental estimates.⁵

Our findings have potential implications for the design of progressive tax reforms. In a nutshell, the government must factor in the tax compliance responses to forecast the true effects of a tax policy (Saez and Zucman, 2023). Once behavioral responses are factored in, tax progressivity may not increase as much as intended, and a reform that was supposed to be revenue-neutral may not be. Our counterfactual analysis evaluates a hypothetical tax reform that reduces the tax liability for poor households by 39.1% and increases it for rich households by 17%. In the absence of behavioral responses, the reform is revenue neutral

⁵For more details about the design and the results of the forecast survey, see Appendix Section H.

and increases the gap in tax rates between rich and poor households (a simple measure of tax progressivity) by 0.5 percentage points. To compare the outcomes with versus without behavioral responses, we take a sufficient statistics approach (Chetty, 2009): we show that the counterfactual analysis depends on the four behavioral elasticities estimated with the quasi-experimental and experimental approaches. Relative to the counterfactual with no behavioral responses, the real effects of the reform are significantly different, with an effective tax progressivity that is 27.6% lower and tax revenues that are 3.8% lower.

Our paper relates and contributes to multiple strands of literature. First, we contribute to the literature on tax progressivity. The topic of tax progressivity received renewed attention. This interest is growing more generally, but also in the context of property taxes, which tend to have flat schedules, that is, proportional to value, contrasting sharply with modern progressive income and wealth taxes (Chancel et al., 2022; Dray et al., 2023). Recent studies on progressivity document its evolution over time and its connection to income inequality (Piketty and Saez, 2007). Other work has examined the optimal level of progressivity, balancing the efficiency-equity trade-off (Heathcote et al., 2017). There have also been substantial advances in our understanding of perceptions and support of progressivity and their implications (Kuziemko et al., 2015; Ballard-Rosa et al., 2017; Stantcheva, 2021; Hoy, 2025). These contributions, while invaluable, have been primarily descriptive, theoretical, or based on stated preferences. Our work builds on these previous findings and contributes with causal estimates of the effects of a real-world progressive tax reform on actual tax compliance in the context of an at-scale field experiment.

Second, our study contributes to the broad literature on social preferences, tax preferences, and preferences for redistribution. Our results on the cross-rate effects of progressive reform relate to recent empirical advances in tax morale and tax compliance (Luttmer and Singhal, 2014; Slemrod, 2019). Previous work has largely relied on survey data to study these preferences. Stantcheva (2021) used detailed survey experiments to show that educational videos on the redistributive effects of tax policy increase support for progressive taxes, while efficiency-focused videos have no effect. Similarly, Hoy (2025) conducted multicountry survey experiments finding that tax morale increases when individuals learn about the progressive nature of their tax systems. Although valuable, these studies rely on stated preferences, which may be subject to social desirability and other biases: rich individuals might express support for higher taxes but behave differently when stakes are real.⁶ We complement this literature by examining support for progressive taxation through revealed preferences in a natural, high-stakes context. Our results highlight important discrepancies between stated and revealed preferences: while middle-income and rich households report higher perceived fairness after learning about the progressive reform, this does not translate into increased compliance. In contrast, poor households show

⁶For more discussion of the generalizability of data on social preferences, see Epper et al. (2024).

consistency between their survey responses and behavior—both their fairness perceptions and tax compliance increase upon learning about the reform. Our approach builds on [Besley et al. \(2023\)](#) who studied how the UK’s regressive poll tax increased tax evasion, by identifying a new tax morale channel—preferences for progressivity—and estimating both own-rate and cross-rate effects in a unified setting.⁷

The rest of the paper proceeds as follows. Section 2 describes the institutional context and data. Section 3 discusses the own-rate effects, while Section 4 discusses the cross-rate effects. Section 5 presents the counterfactual analysis. The last section concludes.

2 Institutional Context and Data

2.1 Local Property Tax

Our study focuses on Tres de Febrero, a major urban municipality in the Greater Buenos Aires metropolitan area, Argentina, which levies a local property tax known as *Tasa por Servicios Generales* (TSG), related to the provision of services such as street lighting, urban sanitation and maintenance.⁸ This type of tax is common across all Argentine municipalities and serves as the primary source of local revenue. Tres de Febrero is no exception, with the TSG accounting for 20% of its total resources and 45% of its own-source revenue in 2021, which is consistent with the share of real estate property taxes for state and local tax revenues in the United States ([Saez and Zucman, 2019b](#)).

This tax has two components: a variable part calculated by applying a tax rate to the property’s assessed value (with rates ranging from 0.32% to 2.48% across eight property categories), and fixed charges for specific services like security and health (Appendix Section A.1 provides full details on the TSG composition). Property assessments are based on the provincial tax authority’s cadastre, but these values differ systematically from market values and are only used for reference purposes.⁹

Before the 2022-24 progressive reforms, the tax structure was mildly regressive due

⁷[Nathan et al. \(2024\)](#) provide related evidence that fairness considerations can influence tax compliance through tax appeals. However, their study focuses on reciprocal fairness—a concept distinct from the type of fairness we examine. They show suggestive evidence that households are more willing to pay taxes when they believe others are contributing their fair share. Similar findings on reciprocal fairness appear in [Hallsworth et al. \(2017\)](#) and [Del Carpio \(2014\)](#).

⁸As of 2022, Tres de Febrero had approximately 365,000 residents and 115,000 households, representing 1.2% of Argentina’s total population.

⁹The relatively high statutory rates are due to a chronic issue of under-valuation of properties in administrative records, which is “solved” by setting high statutory rates. Our best estimate is that properties’ administratively assessed values are about 29.54% of their market values in Tres de Febrero, which results in a effective tax rate of about 0.9%, at about the middle point of the State property taxes range in the United States, and similar to that of Missouri (0.82%) and Maryland (0.96%) in 2022 ([Tax Foundation, 2024](#)).

to the fixed charges comprising a larger share of the total tax burden for lower-value properties. For properties in the bottom decile, the fixed portion represented 70% of the tax bill while the variable portion was 30%. Conversely, for the top decile, the fixed component accounted for about 15%, and the variable portion 85%. In 2022, the average annual tax rate was 3.01% of the assessed property value, with properties in the 25th percentile facing a rate of 2.74% compared to 1.83% for those in the 75th percentile.

2.2 The Progressive Tax Reform

For political and equity reasons (mainly negotiations at the local council level with the center-left opposition), the municipal government implemented a series of progressive reforms to the TSG in three stages between 2022 and 2024. The first change in 2022 introduced a 10% surtax on properties valued at 1.5 million Argentine Pesos (hereinafter, ARS) and above, though this was not widely communicated.¹⁰ In 2023, the municipality implemented a more substantial reform, applying a 30% discount to the variable component for properties valued at ARS 750,000 or less, while maintaining a 10% surcharge for those above ARS 1.5 million. The third stage in 2024 introduced progressivity to the fixed component of the tax, with discounts for low-value properties and surcharges for high-value ones (Appendix Section A.2 provides full details of these reforms). These last two reforms were also introduced and approved rapidly, and they were not advertised beyond the experiment.¹¹

These reforms created three distinct groups with distinct tax changes and sharp and binding boundaries at fixed levels of administrative valuation: Properties at the bottom 40% of the value distribution (poor households) received a reduction in their median tax rate of about 13.62%, those at the top 22% (rich households) experienced an increase in their median rate of approximately 7.44%, and those in the middle of the distribution saw no change in their tax liabilities.

The reforms established sharp discontinuities in tax rates at these thresholds, providing an opportunity to assess the *own-rate* effects of tax changes (see Section 3). In addition, we used the timing of these reforms to conduct an information randomized controlled trial in which we experimentally varied the awareness about the progressivity of the reform to the three groups of taxpayers, to estimate the *cross-rate* effects of these tax changes, as explained in Section 4.

¹⁰Properties valued above ARS 3 million faced additional increases in 2023 and 2024, though we exclude this group from the analysis since it represents less than 6% and we do not have enough statistical power in that range. All of our results remain virtually unchanged when including this group.

¹¹The 2023 tax reform was discussed and decided at a city hall council meeting on December 5, 2022, and published in the annual tax bill on December 13, 2022 ([link](#)).

2.3 Tax Data, Compliance, and Facts

Our analysis relies on monthly administrative data from the municipality’s tax records from January 2018 through February 2024. The primary database consists of monthly property tax payment records constructed from the monthly bills issued to account holders. The unit of observation is an “account,” which coincides with a property unit. The data set contains the following billing details: account number (unique property identifier), address, name of locality (neighborhood), year and month of the bill (12 monthly bills), the monthly due amount (in pesos), a payment indicator, due date, date of payment, days overdue, means of payment (cash or electronic), type of account (residential, retailer, manufacturer), and linear front meters of the lot/property.

In developed countries, compliance rates are typically high and tax delinquency is often not a substantial margin of interest. In developing countries, however, several factors limit the ability of the local tax agency to enforce property taxes. Consequently, there is potential for tax reform to significantly impact tax compliance. For example, as reported in [Nathan et al. \(2024\)](#), in studies conducted in Argentina, Brazil, and Haiti, over 50% of households failed to pay their property taxes on time; in contrast, only 0.42% of households in Dallas County, Texas, failed to pay their property taxes on time. In our context, the delinquency rate of 2022 was still high at about 63% for the TSG.

In Tres de Febrero, property taxes are typically paid on a monthly schedule, with each installment due by the 15th of the respective month. Our main outcome of interest is a binary variable that indicates whether the household made a tax payment at most three months after the due date.¹² This measure has been used in other studies of property tax enforcement ([Del Carpio, 2014](#); [Castro and Scartascini, 2015](#); [Carrillo et al., 2021](#); [Cruces et al., 2024](#)), and the results are robust to alternative definitions of the outcome variable. In the year before the reform (2022), the average compliance rate was 47%. There is significant variation in compliance across households. A large fraction of 42.6% of taxpayers did not make any timely payments in the year – indeed, it is common for some households to go years without paying any property taxes bills at all, timely or not. At the other end, 26.3% of households complied with their taxes every month of the year. The remaining households made timely payments in some but not all months of the year. Finally, we use a variation of this outcome of interest for our regression discontinuity estimates of the *own-rate* effects, as we discuss in detail in the next Section.

¹²For October, November, and December 2022, we restricted this variable by not allowing payments of a pre-treatment installment to be made in the post-treatment period. In cases where such payments were made after the post-treatment period, their value was set to zero.

3 The Impact of Changes on One’s Own-Rate

3.1 Research Design

Our research design exploits two key sources of identification to measure the *own-rate* effects of the progressive tax reform. First, as described in the previous section, the treatment assignment is based on a two-cutoff cumulative design (Cattaneo et al., 2016). Specifically, the reform created two sharp discontinuities in tax rates based on exogenously set property values: households with values below a predetermined threshold received a tax cut while those above did not, and similarly, households above a higher threshold experienced a tax increase while those below it did not. Second, we leverage the sudden and unannounced implementation of the reform, which allows us to compare outcomes before and after its introduction.

A distinctive feature of our approach is that we apply a regression discontinuity framework to study changes rather than levels. In what we label a first stage, we analyze how tax rates changed from one year to the next around these thresholds, to establish that the changes induced by the reform were indeed binding. In the subsequent reduced form analysis, we study how compliance rates changed around those thresholds. To fit this framework, we adapt our outcome of interest by defining it as the change in the proportion of bills paid within three months after the due date between the pre-reform and the post-reform semesters.¹³

Since it leverages the discontinuity around the threshold as well as the changes around the implementation date, our research design is reminiscent of a difference-in-discontinuity approach. This approach is critical for analyzing the 2024 reform replication, which utilized the same thresholds as the 2023 reform. By examining changes rather than levels, we can identify the incremental effect of each reform separately, rather than the cumulative effects of the previous reforms. Moreover, there may have been other past reforms that used the same or similar thresholds. The use of year-over-year changes as dependent variables enables us to focus precisely on the shock from the specific reform year – this is also crucial for the falsification test that uses a fake implementation date. This methodological choice thus strengthens both the internal validity of our main results and our ability to assess the robustness of the findings through replication and falsification.

The RD parameters are estimated using nonparametric local-linear regression with a mean-squared error optimal bandwidth (Calonico et al., 2014). We conduct this regression discontinuity analysis for both the six-month change in the effective average tax rate $R(\tau)$ (our “first stage”) and for the change in the compliance rate by semester before and after

¹³For instance, if an individual paid four bills in the last semester of 2022, and then paid five bills in the first semester of 2023, her outcome is computed as 0.167 p.p.

each reform (our “reduced form”).

3.2 Own-Rate: Main Results

Figure 1 summarizes our main *own-rate* results. Panels (a) and (b) present the change in tax rates around the poor/middle and middle/rich thresholds, respectively. They both depict visible, large and statistically significant effects on the tax rates owed by poor and rich households around the relevant thresholds, comparing the year of the reform (2023) with the previous year. The progressive reform induced a reduction in poor households’ tax rates by about 0.61 percentage points compared to middle-valuation households near the threshold. Conversely, rich households’ tax rates at the threshold increased by about 0.25 percentage points compared to middle-valuation households near the threshold.

The reduced form estimates of the reform on compliance are presented in Panels (c) and (d) of Figure 1. While not as sharp as the jumps in effective tax rates, the estimation results document substantial and significant changes in tax compliance for both poor and rich households. The impact on poor households is large, with an increase of 2.74 p.p in the proportion of bills paid from one semester to the next, while for rich households we find a reduction of 2.29 p.p in the same outcome. In a nutshell, poor households experienced a fall in their tax rates, and responded by increasing their tax compliance, whereas rich households saw an increase in their rates and reduced their compliance.

Our RDD results are robust to a series of alternative estimation choices and tests. The results in Figure 2 are based on the same specification, but use placebo discontinuities of the running variable (500K and 2M instead of 750K and 1.5M) for both the first and the reduced form. Figure 3, in turn, presents the results of placebo dates for the effect on tax compliance of a non-existing reform in mid-2023. Reassuringly, the results in both figures indicate that our main estimates are robust to these placebos. Appendix Figure C.1 shows the result of the manipulation tests around the cutoffs, indicating that we cannot reject the null hypothesis of no discontinuity of the density at the cutoff. Finally, Appendix Table B.1 presents several robustness checks to our preferred specification. We find similar patterns when we add additional pre-reform controls, using an alternative bandwidth, or when using a different definition of the outcome (for instance, including the possibility of paying 6 months after the due date instead of only 3).

3.3 Own-Rate Elasticities

To estimate the *own-rate* effects, we calculate two elasticities based on these regression discontinuity design estimates. $\beta^{own-rate}$ represents the effect from the reduced form of the RDD, $\alpha^{own-rate}$ denotes the effect from the first stage, $C^{own-rate}$ denotes the compliance of

the middle group, just above/below the cutoff, and $\tau^{own-rate}$ represents the tax rate of the middle group, just above/below the cutoff. Both elasticities are expressed in terms of six-month to six-month changes, comparing periods before and after the reform. Specifically, the percentage change in compliance is derived from the reduced form of the RDD results, whereas the percentage change in tax rates is based on the first stage of the RDD results. We establish baseline compliance and tax rates using 2023 levels for the middle group – that is, focusing on middle-value properties situated just above (for poor households) and just below the cutoff (for rich households). We compute confidence intervals at 90% significance levels for these elasticities using 5,000 bootstrap iterations.

$$\varepsilon_{poor}^{own-rate} = \frac{\frac{\beta_{poor}^{own-rate}}{C_{poor}^{own-rate}}}{\frac{\alpha_{poor}^{own-rate}}{\tau_{poor}^{own-rate}}} = \frac{\frac{-2.736}{60.1}}{\frac{0.605}{2.222}} \approx \frac{-0.167}{[-0.256; -0.068]} \quad (1)$$

$$\varepsilon_{rich}^{own-rate} = \frac{\frac{\beta_{rich}^{own-rate}}{C_{rich}^{own-rate}}}{\frac{\alpha_{rich}^{own-rate}}{\tau_{rich}^{own-rate}}} = \frac{\frac{-2.290}{58.3}}{\frac{0.245}{2.239}} \approx \frac{-0.359}{[-0.661; -0.009]} \quad (2)$$

We estimate a significant *own-rate* elasticity for poor households of -0.17, which means that for each 1% decrease in the tax rate of the poor households, they increase their compliance rate by 0.17%. The *own-rate* elasticity for rich households, in turn, is -0.36, which implies that for each 1% increase in the tax rate of the rich households, they reduce their compliance rate by 0.36%. To test whether the difference in elasticities between the two groups is statistically significant, we conducted 5,000 bootstrap iterations, which revealed no significant difference (p-value = 0.309).

Summing up, our estimations indicate that both poor and rich households exhibit substantial changes in their tax compliance behavior when facing a change in their tax rates: compliance falls as taxes increase, and vice versa. The rich seem to respond more strongly than the poor, although these differences do not seem to be highly significant. While we have so far concentrated in the direct effects of each group’s own rate changes on their compliance, in the following section we establish whether there is an indirect effect of changes in other groups’ tax rates on the own group compliance.

4 The Impact of Changes in Others' Rates: The Cross-Rate Effects

4.1 Experimental Design

The RD approach allowed us to capture the direct or *own-rate* effects of the reform. Our main research question, however, is how changes in others' tax rates in the context of a progressive tax reform affect one's own compliance and perceptions of fairness. We designed an at-scale information provision randomized-controlled trial (consisting of about 90,200 letters) to identify these *cross-rate* effects.

The experiment focused on the progressive nature of the property tax reform that decreased taxes for lower-valued properties and increased them for higher-valued ones. This setup provided a unique opportunity to study the effects of progressive taxation in a real-world environment, allowing us to observe and analyze actual taxpayer behavior in response to tax policy changes.

Property tax (TSG) bills are mailed monthly to all taxpayers. These letters convey personalized information to each account holder, such as the property's tax valuation, the payment structure, a monthly breakdown of the tax amount due, their expiration dates, and the payment options. This level of detail ensured that recipients had a clear understanding of their financial obligations. Section D shows an example of a letter sent to the households.

Our information provision RCT leveraged the January 2023 letters to provide details about the changes induced by the reform in taxpayers' obligations. All participating households received their usual tax bill, but the back of this bill also displayed information about the tax reform (and other government measures). All letters included information about the implementation of a tax reform in very broad terms, and emphasizing its aim to enhance the tax system's equity, but without specifying how the reform affected *others'* rates. Importantly, all letters specified whether the recipient's *own* tax rate was reduced, increased, or not impacted by the reform.

The letter mailed to the control group consisted of the tax bill on the front and, at the back, the limited information about the reform and the detailed information on how it affected the recipient's tax rate. The key aspect of our research design is that it experimentally varied the awareness about the tax reform's progressivity by providing a detailed description of the reform to a randomly selected subset of taxpayers. The letters mailed to this treatment group contained exactly the same baseline information as in control group letters, but they also included a key element designed to convey the progressive nature of the reform's—informing not only how it affected the recipient's tax

rate, but also how it affected that of other groups.¹⁴

The key piece of our awareness treatment was the provision of an infographic which highlighted the reform’s progressive nature. The infographic and its accompanying text in the treatment group’s letter were the result of a collaboration with policy communication professionals. It was designed to be visually engaging and informative, and it presented a clear, concise depiction of how the reform impacted the tax rate of different property valuation brackets (the poor, the middle group and the rich) besides the recipient’s own rate in a way that was easy to understand for a non-specialized public.

Since the reform was not widely discussed nor advertised and it was approved and implemented in a relatively short period of time, most taxpayers were probably not aware of it nor of its details. Only about 15.4% of our survey sample did know about the reform (see the discussion of Figure 5 in the next subsection). We expect our information treatment (by means of this visual representation) to have induced not only increased awareness about the reform but also genuine learning about it among a relatively dis-informed public. At the very least, our treatment should have boosted the reform’s salience among those who already knew about it, although this group seems to be only a minority of taxpayers. In any case, because of increased awareness, learning or salience, we expected that this detailed explanation of the reform’s progressive nature would lead to play a significant role in shaping households’ perceptions and attitudes towards the reform, and potentially induce changing patterns in tax compliance and perceptions of tax fairness with respect to the control group.

Because taxpayers were impacted differently by the reform, we stratified our experiment across the three property valuation groups (the poor, middle, and rich groups, as defined above). The comparison of outcomes between recipients of treatment and control letters within each group allows us to capture the impact of changes in others’ tax rates, i.e., *cross-rate* effects, on subsequent tax compliance, and perceptions of fairness.

We provide a series of illustrations of our treatment and control messages. panel (a) of Figure 4 presents an example of the message received by those in the control group in the poor bracket (similar info graphics were included in the letters for rich and middle-valuation taxpayers).¹⁵ This messaging was carefully calibrated to acknowledge the reform’s intent to change the tax burden for the taxpayer without emphasizing its progressive character nor the changes in other groups. The treatment infographic was more explicit and detailed in its depiction of the reform’s progressivity (panel (b) of Figure

¹⁴The inclusion of information about changes in the taxpayer’s own tax rate in both control and treatment letters is a critical aspect of our research design since it allows us to disentangle the effect of the increased awareness about the reform’s progressivity from the direct financial impact of the tax change on the taxpayer’s behavior.

¹⁵The control and treatment messages for middle and high property valuation groups are presented in Appendix Figures D.1 and D.2.

4 presents the message included in the letter for the poor group). It clearly outlined the tax implications for different property valuation groups, indicating which brackets would experience an increase, a fall, or no change in their tax rates, and the magnitude of these changes in the rates.¹⁶ Finally, a last group of taxpayers did not receive any letter. This group allows us to evaluate whether taxpayers typically open correspondence from the government, which is a prerequisite for our treatment to have a meaningful effect, and to compare this basic effect to those in the literature (see the systematic review by [Antinyan and Asatryan, 2024](#)).

4.2 Cross-Rate: Main Results

We implement the research design described above by comparing treated and control taxpayers in the three groups separately. We draw our outcomes on perceptions and awareness from the post-treatment survey, and on tax compliance from the administrative tax records, as described in Section 2.3.

We first document the effect of our intervention on citizens’ awareness of the reform – we asked respondents in our survey whether they knew about this measure or not.¹⁷ Only about 15.4% of respondents in the control group indicate that they were aware of it with relatively similar levels of knowledge among taxpayers in the three property valuation groups, as depicted in panel (a) of Figure 5. This is consistent with our prior that knowledge about the reform was low among taxpayers given the relatively low level of discussion about it in the public sphere. Importantly, panel (b) reveals a consistent treatment effect across all valuation brackets. Our information treatment significantly increased awareness of the reform in the three groups by about 7.9 percentage points on average.

The back of the tax bill letter also included information about two additional policies the Municipality had implemented – soft loans for businesses, and paperwork simplification for new businesses. The survey also asked respondents about awareness of these policies. Unlike our the information about the progressive nature of the reform, which was only present in treatment letters, both letter types included information about these two policies. We use the question about awareness with business loans as a falsification test of our information treatment (the results are very similar with awareness about the

¹⁶The treatment informed taxpayers that the *variable* component of the tax rate for the poor was reduced by 30% and that this component was increased by 15% for the high property valuation group. Because of the TSG fixed component (see Appendix Section A.1 for details), the final impact on the average effective tax rate was an average reduction of 9.1% for the poor group and an average increase of 7.8% for the rich.

¹⁷Because survey respondents received additional information and questions related to the reform, we drop them from the main experimental sample for tax compliance outcomes. Our results remain qualitatively similar if we keep this group in the sample.

paperwork simplification policy). Panel (c) of Figure 5 indicates a relatively low and evenly distributed level of knowledge about this policy among taxpayers, with only 12.5% stating that they knew about it. We can test the specificity of our intervention by evaluating whether the information about the reform had some effect on awareness about this different policy, which would indicate the presence of social desirability bias or other sources of spurious results. Reassuringly, the results of the falsification test in Panel (d) of Figure 5 indicates that this was not the case. There are no treatment effects of information about the reform on participants’ knowledge of this unrelated government program – the overall coefficient and those for the three groups are virtually zero. This absence of an effect on an unrelated policy supports our conclusion that the observed increase in tax reform awareness and any subsequent changes in behavior and perceptions can be directly attributed to our targeted intervention, rather than to a general increase in policy awareness or other spurious effects.

The results from the survey allow us to study the impact of our intervention on our first outcome of interest, taxpayers’ fairness perceptions about the reform. We asked respondents whether they thought the local tax system was fair– more specifically, we asked how fair they considered the distribution of the municipal tax rates between richer and poorer people, on a scale from 0 to 10 (in the context of a survey about the TSG). Panel (a) of Figure 6 indicates relatively high perceived levels of fairness, with an average perception of about 6.5 and relatively evenly distributed between the three groups. Interestingly, our information provision treatment increased this perception across the board, as depicted in panel (b) of Figure 6. There is a positive and significant treatment effect on fairness perceptions across all valuation brackets of about 0.52 points in the 0-10 scale, with fairly similar effects for the poor, middle and rich groups.¹⁸ However, these effects may only be due to priming or to a social desirability bias, with respondents receiving information about a progressive reform simply stating that positive views in general when asked. We included a further question in the survey as a falsification test. We included the General Social Survey on agreement with the role of government in reducing the gap between the rich and the poor, which allegedly represents deeply rooted general preferences and not something about a specific policy. Reassuringly, the results in Panel (d) of Figure 6 indicate that our information treatment did not affect respondents’ redistribution preferences. This null effect on preferences supports our interpretation that our information treatment truly increased the perception of fairness of the tax system and does not represent some spurious result.

Our analysis now turns to the other core finding of our experimental setup: the *cross-rate* effects of the reform on actual taxpayer behavior. Panel (a) of Figure 7 presents the

¹⁸These findings are aligned with previous work that has documented a general preference for progressive tax systems in both laboratory and survey settings (e.g., Durante et al. 2014; Stantcheva 2021; Hoy 2025)

estimated treatment effects by month on our primary outcome measure – the probability of paying the tax bill no later than three months after the due date (see Section 2.3 for details) – between treatment and control taxpayers separately for the poor, middle and rich groups. The pre-reform levels and trends reflect the random assignment of the treatment within each of the three groups: compliance is balanced between treatment and controls, with small and mostly non-significant differences between the two for any of the valuation brackets. Interestingly, there is a clear positive and significant effect of the treatment on compliance for the poor group, as witnessed by the six monthly coefficients. For the middle group the differences are smaller and not statistically significant at conventional levels. Finally, we also document a slightly negative effect on rich households in some months.¹⁹ To provide a comprehensive view of our findings, we aggregated the monthly data into three-month periods, before and after the treatment, while maintaining the distinction between valuation brackets. Panel (b) of Figure 7 presents the ‘After Letter’ coefficients of these pooled results, which align consistently with our previous observations, and the ‘Before Letter’ coefficients, which show no difference between groups. The pooled quarter estimate indicates a positive effect of 0.952 percentage points of the treatment on tax compliance (significant at the 5% level) for taxpayers in the poor group. In contrast, middle-value individuals show no significant change in payment behavior. Rich individuals display a negative pooled estimate of -0.239, albeit not significant at conventional levels. Appendix Figure E.1, in turn, presents the raw differences before and after the intervention between treatment and controls for the three groups. These additional results confirm the pattern in Panel (a) in Figure 7. Pre-intervention, treated and controls taxpayers exhibit comparable levels of compliance. Post-intervention, poor households in the treatment group exhibit a noticeable increase in their probability of making a tax payment, while rich households exhibit a slight decline compared to their control counterparts.

Table B.2 in the Appendix presents several robustness checks. For instance, we find similar patterns when we exclude from the regression control variables not related to pre-treatment payments. We also find similar results when using an alternative sample of subjects (adding the survey sample), when using an alternative definition of the outcome (allowing payments 6 months after the due date instead of only 3), or when excluding outliers in property valuations (below the 5th and above the 95th valuation percentiles).

These results are the more notable when considering that the actual treatment was the inclusion of an infographic and some text at the back of a local tax bill. A potential concern could arise if the letter open rate (or reading rate) was too low, which would bias our results downwards. If taxpayers do not open the letters at all, however, they

¹⁹The patterns seen in this Figure align with existing literature on correspondence studies and suggests that salience may play a role in the observed outcomes. As the message recedes from taxpayers’ immediate memory, its impact appears to wane, a phenomenon consistent with the patterns discussed by Bergolo et al. (2023) in an information provision tax compliance experiment.

would not be exposed to our information treatment, and our results could be spurious. To address this issue, we included an additional simple experiment in our design, in which we compared the effect of receiving a control letter to receiving no letter at all (a small group of taxpayers was assigned to this group). Appendix Figures F.1 and F.2 in the Appendix present these results. Figure F.1 indicates that payment patterns were similar between groups that received a letter and those that did not before the intervention. However, after the letters were delivered, taxpayers in the control group who received a letter consistently and substantially increased their payments when compared to those who received no letter. This pattern was similar across poor, middle, and rich households. Panel (b) of Figure F.2 confirms the results when pooling the outcomes in three-months brackets. The average difference of about 5 percentage point is in line with a recent systematic review of simple experiments of this type (Antinyan and Asatryan, 2024).²⁰ Overall, the results from this auxiliary experiment suggests that probably a significant share of taxpayers opened the letters that we sent them.

However, opening the letters does not necessarily mean reading them carefully. Our treatment was intentionally subtle, with the message appearing only on the back page of the tax bill letter. The effect of sending the information provision letters (0.952 p.p. for the poor group, for instance) should be interpreted as an intention-to-treat (ITT) coefficient. The Average Treatment Effect on the Treated (ATET) can be considered a factor of the ATE, which is affected by the rate of non-compliance. For a taxpayer to be actually treated, the letters needed to arrive on time, be opened, read front and back, and the recipient must process and understand the message. The recent literature on similar experiments estimates these effects for the United States, and finds non-compliance adjustment factors between 6.5 and 7.5 (Perez-Truglia and Cruces, 2017; Bottan and Perez-Truglia, 2020; Gerber et al., 2020; Nathan et al., 2020). In our context, we can use the results from Figure 5 (Panels (a) and (b)) to estimate the treatment effect on the probability of knowing the reform’s progressive nature as a proxy for compliance with the treatment. If 20% of poor households in the control group were aware of the reform and 100% in the treatment group were exposed to the information, we would expect awareness to increase by 80 p.p due to our treatment (from 20% to 100%). However, we identify an effect of approximately 10 p.p (Panel b), implying a reading rate of 12.5% (10/80), or a non-compliance adjustment factor of 8. This adjustment factor suggests that the treatment ATET on the poor was 8 p.p., which we consider an upper bound and interpret cautiously. The reading rates we estimate from the survey have several limitations, including the fact that the specific question used to measure awareness included different

²⁰It should be noted that experiments of this type find large effects because they compare the receipt of a letter to no letter at all, which includes a large reminder effect over and above the content of the actual letter. Our information provision experiment’s results are substantially smaller because they rely on the comparison of subtle variations between two types of letters, which nets out the basic reminder effect.

options and reforms, potentially confusing some respondents. Moreover, some taxpayers might have known about the progressive reform but not paid attention to the survey. On the other hand, [Perez-Truglia and Cruces \(2017\)](#) state that, according to the Environmental Protection Agency, nearly 50% of unsolicited mail is discarded unopened, determining a non-compliance adjustment factor of 2, which we will use as our conservative lower bound. Overall, we estimate the effect to be between 2 p.p and 8 p.p, depending on the non-compliance adjustment factor used.²¹

The information provision results on tax compliance, an actual behavior, contrast with those on perceptions documented in panel (b) of Figure 6. First, as panel (a) of the same figure shows using data from the control group, citizens across valuation brackets consider that the TSG is somewhat fair (measured from 0 to 10). Second, as panel (b) shows, using survey self-reported data on perceived tax fairness, we find a significant and large positive treatment effect on perception of fairness (between 0.5 and 0.75 points, in a scale from 0 to 10, with a baseline around 6). This effect is fairly similar across valuation brackets, which suggests that poor and middle but also rich households increased their perception of fairness of the tax system when exposed to our treatment. In addition, as Appendix Figure G.1 suggests, most of the taxpayers surveyed, regardless of their property value, reported being satisfied with a reform that lowers the TSG for the poor and raises it for the rich.

Based on these stated perceptions, we would expect a similar behavioral effect in terms of actual tax payments across the three valuation brackets. On the contrary, the treatment effect on actual tax-paying behavior appears to be mediated by how the reform affects their own tax rates: poor households react by further increasing their tax compliance (over and above the positive *own-rate* effect), while rich households do not seem to react in this dimension (if anything, they reduce their compliance slightly by an amount we cannot detect with our sample size, taking into account that the sample of rich households is less than half as large as the sample of poor households), as illustrated by the results of panel (b) of Figure 7.

²¹Non-compliance is not exclusive of the *cross-rate* effects. The *own-rate* effects (Section 3) could also be plausibly affected by attenuation bias, because of non-compliance. This is because not every taxpayer was necessarily aware of how her own tax changed. However, in the case of *own-rate* estimates, estimating the ATET is considerably more challenging and implies more assumptions. In particular, *own-rates* effect are potentially explained by a mix of a "price" effect (which are not affected by attenuation bias) and tax morale (which are) and thus estimating the scaling-up factor requires an assumption about the relative weight of each of the two effects. We thus prefer not to report scaled-up (ATET) results for the *own-rate* effects.

4.3 Cross-Rate Elasticities

The *cross-rate* effects we estimated allow us to compute two relevant elasticities. These will be based on the pooled results of 3 months. $\beta^{cross-rate}$ represents the effect from the experimental design for each property valuation group, $C^{cross-rate}$ denotes the compliance of each group in 2022, and $\Delta\tau^{cross-rate}$ represents the tax rate change of the *other* property group. Both elasticities are expressed in terms of changes from six months to six months, comparing periods before and after the reform. For simplicity, we characterize the tax changes as a 30% decrease in the tax rate for the poor group and a 15% increase for the rich group, as we conveyed (with respect to the variable component of the TSG) in our sent letters. In computing these *cross-rate* elasticities, we use the percentage rate change of the opposite group in the denominator, reflecting the cross-group nature of the effect. The baseline compliance rates, used as reference points, are derived from the 2023 compliance rates of the control group for each valuation bracket.²²

$$\varepsilon_{poor}^{cross-rate} = \frac{\frac{\beta_{poor}^{cross-rate}}{C_{poor}^{cross-rate}}}{\Delta\tau_{rich}} = \frac{\frac{0.952}{55.2}}{0.15} \approx \frac{0.115}{[0.048;0.181]} \quad (3)$$

$$\varepsilon_{rich}^{cross-rate} = \frac{\frac{\beta_{rich}^{cross-rate}}{C_{rich}^{cross-rate}}}{\Delta\tau_{poor}} = \frac{\frac{-0.239}{54.2}}{-0.3} \approx \frac{0.015}{[-0.029;0.061]} \quad (4)$$

We estimate a significant *cross-rate* elasticity for poor households of 0.115, which means that for each 1% increase in the tax rate of the rich households, poor households increase their compliance rate by 0.115% when they become aware of the change in the other group. We also estimate a *cross-rate* elasticity for rich households of 0.015, although this is not statistically significant at conventional levels. These results are not scaled-up for non-compliance and thus should be interpreted as ITT elasticities. Furthermore, using 5,000 bootstrap iterations, we tested whether the difference in elasticities between poor and rich households is statistically significant, finding that it is, at the 5% significance level (p-value = 0.040).

To put the magnitude of the cross-rate effect in context, we can compare it to the quasi-experimental estimates for the *own-rate* effect: a 1% increase in the tax rate of the rich has the same effect on the compliance of a poor household as lowering the poor household's own tax rate by 0.66% ($\varepsilon_{poor}^{own-rate} \cdot 0.66\% = 0.11\%$).

²²We also calculated confidence intervals at 90% significance using 5,000 bootstrap iterations for these elasticities.

4.4 Further Results: Replication

After the 2023 reform and with evidence that there were positive effects on tax revenue and perceptions, the local government passed a further progressive reform in 2024, which targeted the fixed component of the TSG and resulted in similar coefficients of changes in rates for the poor and rich groups. Besides the detail about the variable or fixed component that saw a change, the messages we sent were very similar between the two years. There were some contextual differences, and we also modified some aspects of the design to establish the robustness of our findings.

First, there were some subtle differences between the two reforms. In contrast to the 2023 experiment, where tax reductions for the poor and increases for the wealthy were applied to the variable component of municipal taxes, tied to property value, in 2024 these adjustments were made to the fixed component of the municipal tax.

Second, the 2024 reform presented an opportunity to refine our approach further. As shown in Appendix J, the treatment letters for the 2024 and 2023 reforms remained identical. However, the 2024 control letter incorporated a subtle enhancement: a small graphical representation (in the same spirit as the treatment group’s info graphic) illustrating by how much the tax rate for the recipient’s taxpayer bracket (poor, middle, or rich) was affected by the reform. This nuanced difference built upon the 2023 control letter, which already included information about the direction of the change in the recipients’ own rates but not the actual magnitude of the change, which was included as part of the info graphic for the treatment group. This replication, thus, makes even more comparable the information about one’s own tax change between the treatment and the control group.

The similarities between the reforms and letters across both years raise the possibility that some individuals in 2024 may have misinterpreted our communications as referring to the previous year’s reform. Moreover, the 2024 reform and its accompanying letters were disseminated during a particularly turbulent economic and political climate, coinciding with one of the highest inflation rates in Argentina’s recent history. These contextual factors could have potentially diluted the anticipated impact of both our treatment and the reform itself.

With these considerations in mind, Appendices L, K and M present our main *own-* and *cross-rate* findings using data from the 2024 reform and the 2024 information provision experiment, mirroring our analysis of the 2023 reform. Consistent with our previous results, we again identify a significant impact across all valuation brackets when comparing recipients of a control letter with those who received no letter (Figures K.1 to K.2). The magnitude of this effect, ranging between 4 and 8 percentage points, remains comparable across both reform cycles.

Moreover, Figures L.1 to L.2 and Figure M.1 illustrate similar patterns in *own-rate* and *cross-rate* effects. We observe a significant (at a 10% significance level) positive cross-rate effect for poor households of approximately 1 percentage point, while middle households show no discernible effects. As in the 2023 experiment, rich households displayed a negative, albeit statistically insignificant, effect of our information provision treatment. Notably, the *own-rate* effects on both the tax amounts owed (first stage) and the compliance rates (reduced form) are more pronounced for poor and rich households in the 2024 reform compared to the 2023 reform.

These results suggest a consistent pattern of behavioral responses to tax information across different economic contexts. The persistence of these patterns, despite the challenging economic environment surrounding the 2024 reform, underscores the robustness of our findings.

5 Counterfactual Analysis

5.1 Counterfactual Setup

In this section, we leverage our results to conduct a counterfactual analysis as a way to illustrate the policy implications of our framework, which incorporates the distinction between *own-* and *cross-rate* effects, and to highlight the potential contribution of the latter to the standard analysis of behavioral responses to tax reforms. We consider the effects of a simplified hypothetical progressive reform that, just like the actual reform that took place in January 2023, reduced the tax rate for properties in the poor group by some proportion $\% \Delta \tau_{poor}$ and increased the tax rate on properties in the rich group by $\% \Delta \tau_{rich}$. We calibrate these factors so that in a world without behavioral responses (i.e., with tax compliance held constant), the reform satisfies the following criteria: (i) it is revenue neutral; (ii) the rich pay an effective tax rate 0.5 p.p. higher than that of the poor. This is achieved by setting a change of $\% \Delta \tau_{poor} = -39.1\%$ and of $\% \Delta \tau_{rich} = +17\%$ to the tax rates of each group. We want to understand the effects of this reform on effective tax progressivity and tax revenues in the presence of behavioral responses, and contrast it to a counterfactual scenario where behavioral responses are not factored in.²³

In this hypothetical reform, we assume that all households within each property value

²³Throughout this section, we use *effective* tax progressivity, which is different from *statutory* tax progressivity. The latter refers to the progression of tax rates as determined by law – essentially, how tax rates increase with property values. Effective tax progressivity, in contrast, is computed as the ratio of actual tax payments to property values along the distribution. This implies that effective progressivity incorporates tax compliance, and thus reflects the difference in taxes actually paid by different property valuation groups after accounting for behavioral responses and enforcement levels. This distinction is critical to understanding the implications of tax reforms, as statutory progressivity does not always translate into effective progressivity because of behavioral responses.

bracket are affected equally, which allows us to concentrate on average effects within each group.²⁴ Let subscript $j \in \{poor, middle, rich\}$ denote poor, middle and rich households.²⁵ Let γ_j be the share of households in each of these three groups. Let superscript s denote the scenarios: $s = 0$ represents the status quo (i.e., before the reform), while $s = A$ is the *actual* post-reform scenario (that is, accounting for all behavioral responses) and $s = C$ is the *counter-factual* post-reform scenario but without behavioral responses. Let $L_j^s > 0$ denote the average tax liability. For example, L_{poor}^0 is the average liability for poor households before the reform. Let $C_j^s \in (0, 1)$ denote the average tax compliance. For instance, C_{poor}^0 is the average probability that a poor household pays its taxes on time before the reform. Lastly, let $T_j^s = L_j^s \cdot C_j^s$ be the amount of taxes paid on time. For this counterfactual analysis, we fix the values $\{\gamma_j, L_j^0, C_j^0\} \forall j$ to be equal to the corresponding averages during the 12 months before to the 2023 reform, from January through December 2022.

The progressive reform can have effects on taxes paid via two channels. The first is the mechanical channel: lowering tax rates for the poor reduces their tax liabilities, whereas increasing tax rates for the rich increases theirs. In the counterfactual scenario with no behavioral responses, this is the only channel at play. In the post-reform scenario with behavioral responses, there is an additional channel operating through compliance: i.e., holding constant the new tax liabilities (L_j^s), the amount of taxes paid (T_j^s) changes because of changes in tax compliance (C_j^s). This second channel can be summarized with four key parameters: the elasticities $\varepsilon_{poor}^{own-rate}$, $\varepsilon_{rich}^{own-rate}$, $\varepsilon_{poor}^{cross-rate}$ and $\varepsilon_{rich}^{cross-rate}$. Following the more conservative estimates discussed in Section 4 above, we adjust the raw *cross-rate* elasticities by a factor of two, in order to scale-up the intention-to-treat effects into treatment effects on the treated. We can then compute the compliance responses to the tax reform as follows:

$$C_{poor}^A = C_{poor}^0 \cdot \left(1 + \varepsilon_{poor}^{own-rate} \cdot \Delta\tau_{poor} + \varepsilon_{poor}^{cross-rate} \cdot \Delta\tau_{rich}\right) \quad (5)$$

$$C_{rich}^A = C_{rich}^0 \cdot \left(1 + \varepsilon_{rich}^{own-rate} \cdot \Delta\tau_{rich} + \varepsilon_{rich}^{cross-rate} \cdot \Delta\tau_{poor}\right) \quad (6)$$

With these compliance rates, we can compute the two main outcomes of interest of the counterfactual exercise. The first outcome is the effective tax progressivity, given by the difference in average taxes paid by rich and poor households in relation to their corresponding average valuations: $P^s = \frac{T_{rich}^s}{\text{AvgValuation}_{rich}} - \frac{T_{poor}^s}{\text{AvgValuation}_{poor}}$. The second outcome is per capita tax revenues, given by the weighted average of taxes paid by each group:

²⁴Because of the fixed and variable components of the TSG (see Appendix Section A.1) there is some heterogeneity on the actual impact of the reform on effective tax rates because it applied only to the variable component—in proportional terms, it affected some households more than others.

²⁵To be consistent with the rest of the analysis in the paper, we exclude the top-6% highest valuations from the sample.

$$R^s = \gamma_{poor} \cdot T_{poor}^s + \gamma_{middle} \cdot T_{middle}^s + \gamma_{rich} \cdot T_{rich}^s.$$

Note that the C_j^A 's are functions of the four behavioral elasticities. We use bootstrap to conduct inference on these magnitudes and account for the estimation error in these four elasticities. Within each bootstrap sample, we estimate the four elasticities of interest and use those values to compute C_j^A following equations (5) and (6). We then use the bootstrap distribution of the statistic to compute the corresponding 90% confidence intervals.

Finally, our framework also allows us to gauge the relative contributions of the *cross-rate* and *own-rate* effects to the overall behavioral response to the tax reform. An advantage of our setup is that it allows us to measure and compare own and social effects within a unified framework and with the same metric. This additional exercise is described in detail in Appendix Section O.

There are some additional implicit assumptions in this exercise that are worth noting. First, to simplify the exposition of the results, we assume that there are no behavioral responses from households in the middle property valuation group.²⁶ The *own-rate* elasticities are estimated by means of a regression discontinuity design and thus correspond to the local average treatment effects for households near the corresponding thresholds. We use these estimates for all taxpayers within each group, implicitly assuming that the effects estimated near the thresholds are good approximations for the average effects within each group.²⁷

5.2 Counterfactual Results

The key results of our counterfactual analysis are summarized in Figure 8 – for more detailed results, see Appendix N. Panel (a) shows the results corresponding to the first outcome of interest: the effective tax progressivity. The key insight from this figure is that a reform initially intended to be progressive may appear less so once behavioral responses are taken into account. The first bar from panel (a) shows that before the reform taxpayers in the rich group faced on average an effective tax rate that was 0.02 p.p. *lower* than that of taxpayers in the poor group—more precisely, the rich paid an effective rate of 0.92% while the poor paid an effective rate of 0.94%. These effective rates thus implied a slightly regressive tax system. The second bar shows that the reform should be expected to make taxes significantly more progressive when ignoring behavioral responses, with taxpayers in the rich group paying an effective tax rate 0.5 p.p higher than that of the poor group

²⁶For the *own-rate* channel, this is true by construction since the tax rates for this group were not affected by the reform and thus their compliance should not change either. The experimental estimates of the *cross-rate* effects for the middle group were small and not statistically significant, and thus we assume a nil behavioral response through this channel.

²⁷For the *cross-rate* elasticity, we are implicitly assuming that there is no treatment heterogeneity with respect to the amount owed. This assumption can be relaxed to obtain more detailed results.

– more precisely, in this scenario the rich pay an effective tax rate that is about twice as high (1.07%) as that of poor (0.57%). The third bar incorporates the behavioral responses from taxpayers in the poor group only. Since taxpayers in this group households react by increasing their compliance, they end up paying more in taxes and consequently this works in the opposite direction of the reform, by narrowing the gap between those in the rich and the poor groups. More precisely, the behavioral responses from those in the poor group reduce the effective tax progressivity by 0.06 p.p when compared to the situation without this group’s behavioral responses. This effect is precisely estimated and highly statistically significant (p-value<0.001). The fourth bar from panel (b) shows our results on tax progressivity when we incorporate the behavioral responses from those in the rich group only. Since taxpayers in this group actually reduce their compliance, total taxes paid by this group fall and thus effective tax progressivity is also reduced. In magnitude, the effect of the behavioral responses from the rich is even larger than the corresponding effect from the poor, reducing the effective tax progressivity by 0.078 p.p (p-value=0.059). The fifth and final bar incorporates the behavioral responses from both groups. Since the two effects go in the same direction, the net effect is a reduction in effective progressivity. In net terms, the behavioral responses lower the effective tax progressivity by 0.138 p.p when compared to the scenario without behavioral responses (p-value=0.002). In other words, due to the behavioral responses, the effective progressivity is 27.6% lower than the scenario without these responses.

Panel (b) of Figure 8 is analogous to panel (a), but focuses on per capita revenues instead of effective tax progressivity. Each bar shows the per-capita revenue under the different scenarios. The first bar from panel (b) shows that before the reform, the per capita revenue was ARS 9,213. The second bar, corresponding to post-reform without behavioral responses ($s = C$), has the same per-capita revenue as in the pre-reform scenario. This is by construction, as we calibrated the reform to be revenue-neutral under no behavioral responses. The third bar shows the per capita revenues when we incorporate behavioral responses from taxpayers in the poor group only.²⁸ Behavioral responses from those in the poor group increase compliance and thus increase per capita revenues. More precisely, this channel generates ARS 113 in additional per capita revenues, an effect that is highly statistically significant (p-value<0.001). In turn, the fourth bar incorporates the behavioral responses from taxpayers in the rich group only. Since those in this group reduce their compliance, the behavioral responses have a negative effect of ARS 347 on tax revenues (p-value=0.059). The fifth and final bar shows the results when allowing for behavioral responses from taxpayers in both groups. Since the two behavioral responses go in opposite directions, the net impact hinges on which of these two offsetting effects predominates. The negative effect on revenues prevails, meaning that even though the

²⁸In practice, this operates by simply taking equation (5) and setting the behavioral elasticities corresponding to taxpayers in the rich group to zero.

reform was designed to be revenue-neutral, behavioral responses may end up reducing tax revenues. However, both the direction and magnitude of this total result should be interpreted cautiously because the total effect is imprecisely estimated.

6 Conclusions

Tax progressivity is a cornerstone of modern fiscal policy, widely adopted to reduce inequality and foster social equity. Although the redistributive effects of progressive taxation are well-documented, empirical evidence on its behavioral impacts—particularly regarding tax compliance—remains limited. This paper fills this gap by examining how progressive taxation shapes compliance behavior, combining quasi-experimental and experimental approaches in the context of a municipal property tax reform in Argentina.

Our analysis reveals asymmetric responses to tax rate changes across income groups: tax reductions for lower-income households significantly increase their compliance rates, while tax increases for high-income households lead to decreased compliance. These effects are amplified by social effects: when informed about tax increases on wealthy households, lower-income taxpayers exhibit further improvements in compliance, concurrent with enhanced perceptions of tax system fairness. These findings demonstrate that behavioral responses to tax progressivity extend beyond direct own-rate effects, suggesting that taxpayer decisions are shaped not only by pecuniary incentives but also by perceptions of distributional equity within the tax system. On the other hand, wealthy and middle-income households also report increased perceptions of fairness under the progressive scheme, but their compliance rates show, if anything, a negative response.

Our findings have important implications for the design of progressive tax reforms. Our counterfactual analysis shows that accounting for behavioral responses significantly reduces both effective progressivity and total revenue compared to projections that ignore these responses. These results underscore that policymakers, particularly in contexts with limited enforcement, must account for heterogeneous compliance responses to achieve their intended distributional and revenue objectives.

From a broader perspective, property taxes represent one of the oldest forms of taxation, and their design has not changed much over time (Chancel et al., 2022; Dray et al., 2023). Despite their crucial role in municipal finance worldwide, with local authorities heavily relying on real estate and land taxes, their structure has received relatively little scrutiny. In most countries, property taxes are either flat (proportional to value) or even regressive due to the prominence of fixed fees, which stands in sharp contrast to the progressive nature of modern income and wealth taxation. However, the recent surge in research on wealth inequality has renewed interest in making property taxes more pro-

gressive. We examine a progressive property tax reform in Argentina, providing novel evidence on both the behavioral responses to such reforms and their broader implications for tax fairness and compliance. Our findings suggest that while progressive property taxation can achieve redistributive goals, policymakers must carefully consider how taxpayers' responses – both in terms of compliance behavior and perceptions of fairness – may affect the reform's ultimate impact on revenue and inequality. As research on wealth inequality and tax fairness evolves, exploring the potential of property tax reforms as a redistributive tool and their long-run efficiency implications, alongside other progressive tax instruments, remains an important avenue for future research.

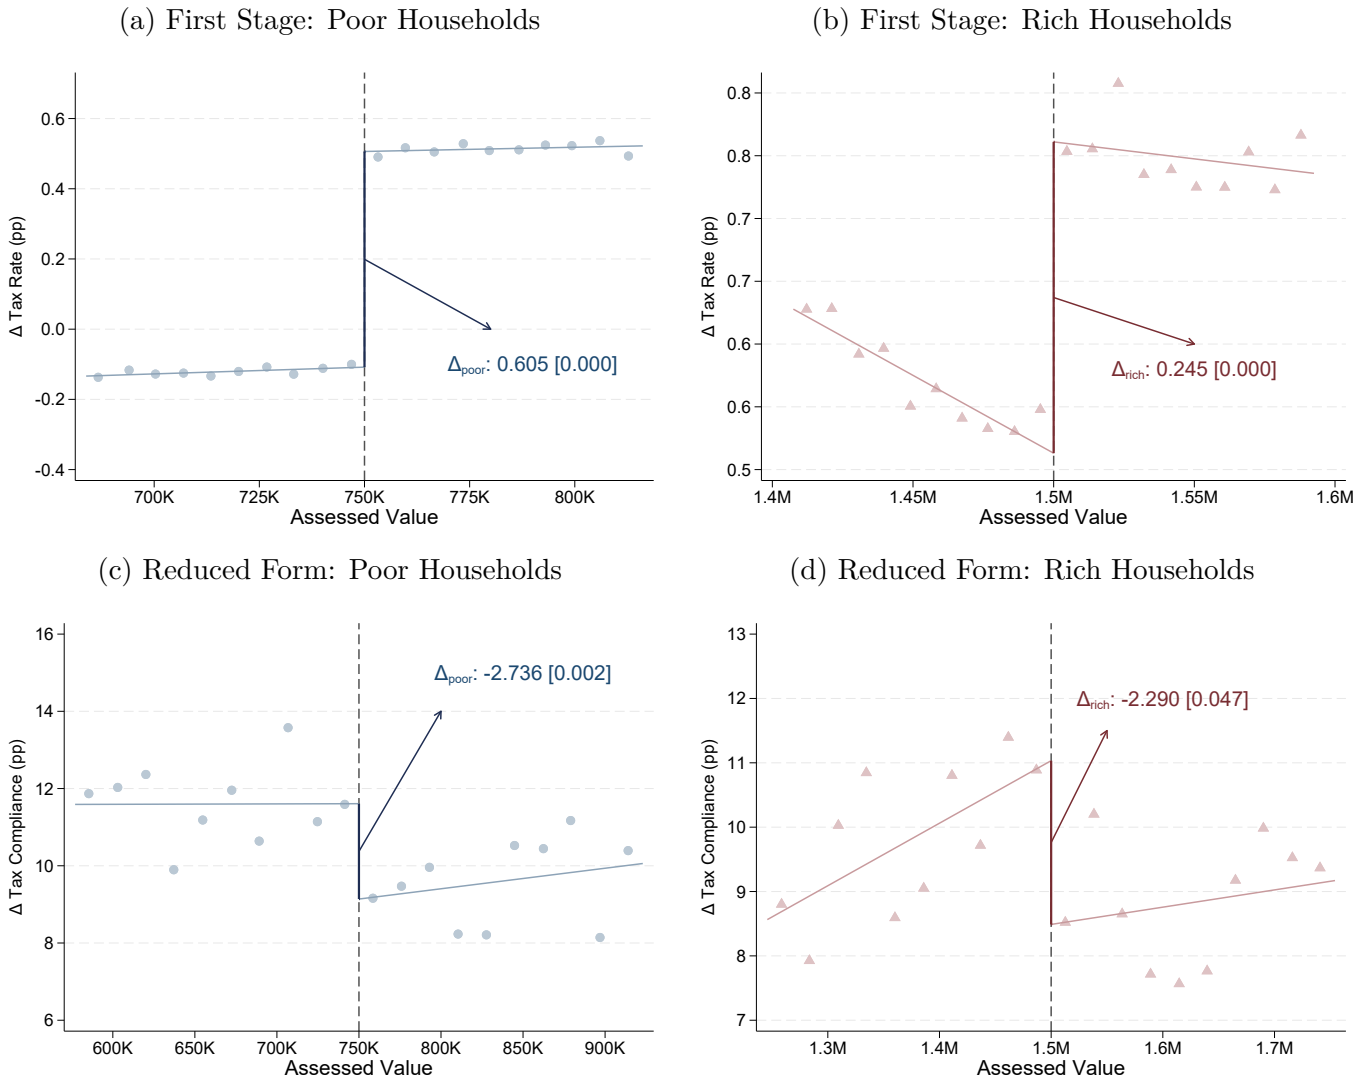
References

- Antinyan, A. and Asatryan, Z. (2024). Nudging for tax compliance: A meta-analysis. *The Economic Journal*, page ueae088.
- Ballard-Rosa, C., Martin, L., and Scheve, K. (2017). The structure of american income tax policy preferences. *The Journal of Politics*, 79(1):1–16.
- Bergeron, A., Tourek, G., and Weigel, J. L. (2024). The state capacity ceiling on tax rates: Evidence from randomized tax abatements in the drc. *Econometrica*, 92(4):1163–1193.
- Bergolo, M., Ceni, R., Cruces, G., Giacobasso, M., and Perez-Truglia, R. (2023). Tax audits as scarecrows: Evidence from a large-scale field experiment. *American Economic Journal: Economic Policy*, 15(1):11053.
- Besley, T., Jensen, A., and Persson, T. (2023). Norms, enforcement, and tax evasion. *Review of Economics and Statistics*, 105(4):998–1007.
- Bottan, N. L. and Perez-Truglia, R. (2020). Betting on the house: Subjective expectations and market choices. *American Economic Journal: Applied Economics*, forthcoming.
- Brockmeyer, A., Estefan, A., Ramírez Arras, K., and Suárez Serrato, J. C. (2023). Taxing property in developing countries: Theory and evidence from mexico. Technical report.
- Calonico, S., Cattaneo, M. D., and Titiunik, R. (2014). Robust nonparametric confidence intervals for regression-discontinuity designs. *Econometrica*, 82(6):2295–2326.
- Carrillo, P. E., Castro, E., and Scartascini, C. (2021). Public good provision and property tax compliance: Evidence from a natural experiment. *Journal of Public Economics*, 198:104422.
- Castro, L. and Scartascini, C. (2015). Tax compliance and enforcement in the pampas evidence from a field experiment. *Journal of Economic Behavior & Organization*, 116:65–82.
- Cattaneo, M. D., Keele, L., Titiunik, R., and Vazquez-Bare, G. (2016). Interpreting regression discontinuity designs with multiple cutoffs. *Journal of Politics*, 78(4):1229–1248.
- CEDLAS (2024). Socio-economic database for latin america and the caribbean (SEDLAC). Joint project with the World Bank’s Poverty and Equity Global Practice.
- Chancel, L., Piketty, T., Saez, E., and Zucman, G. (2022). *World Inequality Report 2022*. Harvard University Press.
- Chetty, R. (2009). Sufficient statistics for welfare analysis: A bridge between structural and reduced-form methods. *Annual Review of Economics*, 1:451–488.
- Christensen, D. and Garfias, F. (2021). The politics of property taxation: Fiscal infrastructure and electoral incentives in brazil. *The Journal of Politics*, 83(4):1399–1416.
- Cruces, G., Tortarolo, D., and Vazquez-Bare, G. (2024). Design of partial population experiments with an application to spillovers in tax compliance. *Review of Economics and Statistics*, forthcoming.
- Del Carpio, L. (2014). Are the neighbors cheating? Evidence from a social norm experiment on property taxes in Peru. *Unpublished Manuscript, Princeton University*.
- Dray, S., Landais, C., and Stantcheva, S. (2023). Wealth and property taxation in the united states. *NBER Working Paper No. 31080*.

- Durante, R., Putterman, L., and Van der Weele, J. (2014). Preferences for redistribution and perception of fairness: An experimental study. *Journal of the European Economic Association*, 12(4):1059–1086.
- Epper, T., Fehr, E., and Senn, J. (2024). Social preferences and redistributive politics. *Review of Economics and Statistics*, forthcoming.
- Fisher-Post, M. and Gethin, A. (2023). Government redistribution and development: Global estimates of tax and transfer progressivity, 1980-2019.
- Gerber, A., Hoffman, M., Morgan, J., and Raymond, C. (2020). One in a million: Field experiments on perceived closeness of the election and voter turnout. *American Economic Journal: Applied Economics*, 12(3):287325.
- Hallsworth, M., List, J. A., Metcalfe, R. D., and Vlaev, I. (2017). The behavioralist as tax collector: Using natural field experiments to enhance tax compliance. *Journal of public economics*, 148:14–31.
- Heathcote, J., Storesletten, K., and Violante, G. L. (2017). Optimal tax progressivity: An analytical framework. *The Quarterly Journal of Economics*, 132(4):1693–1754.
- Hoy, C. (2025). How does progressivity impact tax morale? experimental evidence across developing countries. *Journal of Development Economics*, 172:103398.
- Kuziemko, I., Norton, M. I., Saez, E., and Stantcheva, S. (2015). How elastic are preferences for redistribution? Evidence from randomized survey experiments. *American Economic Review*, 105(4):1478–1508.
- Lustig, N. (2023). *Commitment to equity handbook: Estimating the impact of fiscal policy on inequality and poverty*. Brookings Institution Press.
- Luttmer, E. F. P. and Singhal, M. (2014). Tax Morale. *Journal of Economic Perspectives*, 28(4):149–168.
- Nathan, B., Perez-Truglia, R., and Zentner, A. (2020). My Taxes are Too Darn High: Why Do Households Protest their Taxes? *American Economic Journal: Economic Policy*, forthcoming.
- Nathan, B., Perez-Truglia, R., and Zentner, A. (2024). Paying your fair share: Perceived fairness and tax compliance. *NBER Working Paper No. 32588*.
- OECD (2018). Income Inequality and Poverty in Cities. Technical report, OECD.
- Perez-Truglia, R. and Cruces, G. (2017). Partisan interactions: Evidence from a field experiment in the united states. *Journal of Political Economy*, 125(4):1208–1243.
- Piketty, T. and Saez, E. (2007). How progressive is the us federal tax system? a historical and international perspective. *Journal of Economic Perspectives*, 21(1):3–24.
- Saez, E. and Zucman, G. (2019a). Progressive wealth taxation. *Brookings Papers on Economic Activity*, 2019(2):437–511.
- Saez, E. and Zucman, G. (2019b). *The Triumph of Injustice: How the Rich Dodge Taxes and How to Make Them Pay*. W. W. Norton.
- Saez, E. and Zucman, G. (2023). Distributional tax analysis in theory and practice: Harberger meets diamond-mirrlees. *NBER Working Paper No. 31912*.

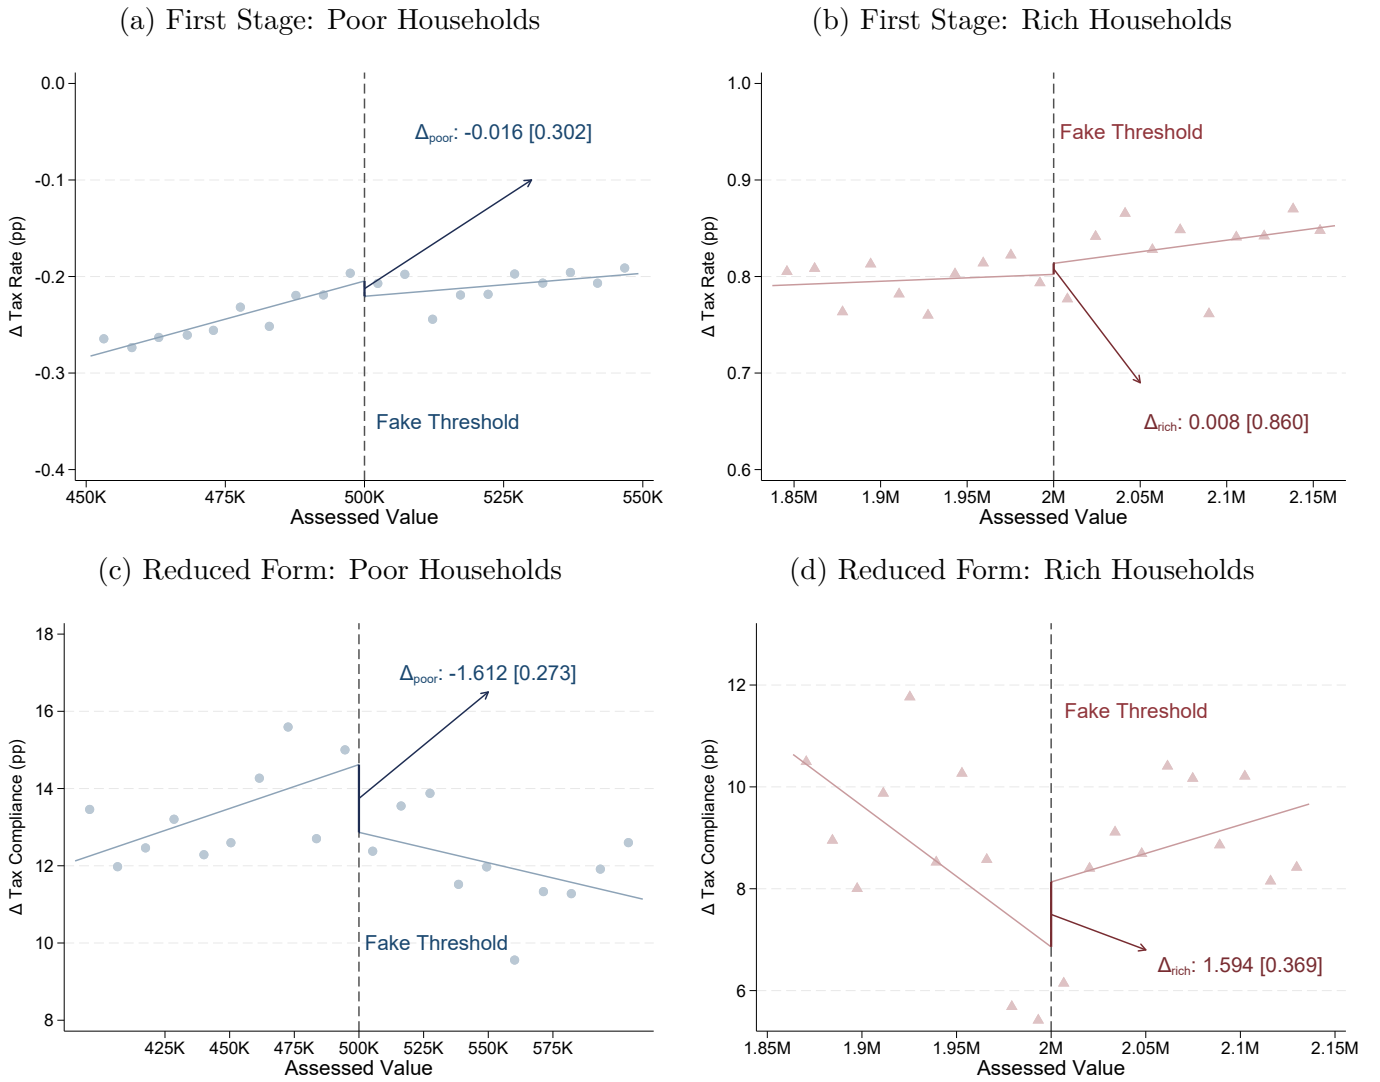
- Slack, E. and Bird, R. M. (2014). The political economy of property tax reform. *OECD Working Paper on Fiscal Federalism No. 18*.
- Slemrod, J. (2019). Tax Compliance and Enforcement. *Journal of Economic Literature*, 57(4):904–954.
- Stantcheva, S. (2021). Understanding tax policy: How do people reason? *The Quarterly Journal of Economics*, 136(4):2309–2369.
- Tarroux, B. (2019). The value of tax progressivity: Evidence from survey experiments. *Journal of Public Economics*, 179:104068.
- Tax Foundation (2024). Property taxes by state & county, 2024. *Tax Foundation Research*. Accessed: 2024-12-10.

Figure 1: Own-Rate RDD Effects on Tax Compliance for 2023



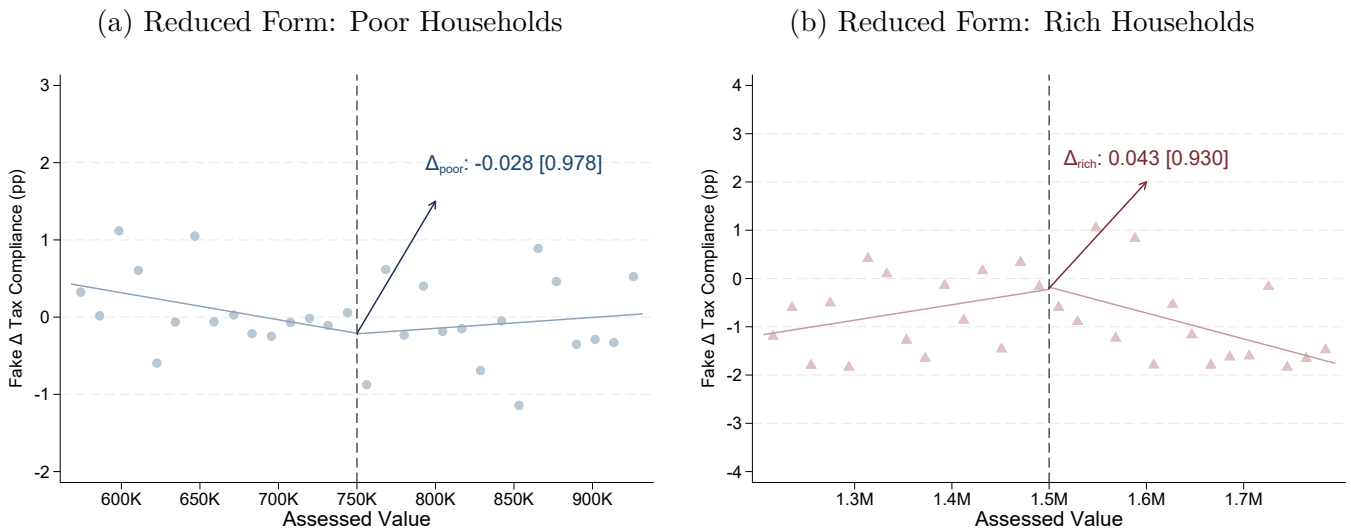
Notes: This figure calculates the *own-rate* effects of the progressive property tax reform in 2023, using a two-cutoff Regression Discontinuity (RD) design. The left figures (in blue) show the results for the poor at the cutoff of 750K. The right figures (in red) show the results for the rich at the cutoff of 1.5m. The running variable corresponds to the properties' cadastral values from 2021. The top panels present the first-stage change in tax rates around these thresholds. The outcome is the household-level change in the tax rate between the last semester of 2022 and the first semester of 2023. The bottom panels present the reduced form effect on tax compliance: how household-level payment rates changed around those thresholds. The outcome is the change in the proportion of monthly bills paid timely between the pre-reform and the post-reform semesters. We define timely payments as bills paid within three months after the due date. Each figure indicates the estimated effect around the cut-off point, where the p-value, taken from a robust bias-corrected inference, is indicated in parentheses. So, the RD point estimates correspond to a change in a time-differenced outcome, reminiscent of a difference-in-discontinuity approach.

Figure 2: RDD Falsification Test using Fake Thresholds for 2023



Notes: This figure calculates the *own-rate* effects using fake tax thresholds at 500K and 2M. Panels (a) and (b) represent the first stage RDD for poor and rich, respectively. These two figures evaluate the increase in the amounts owed between the last semester of 2022 and the first semester of 2023, in relation to the value of the property. Panels (c) and (d) represent the reduced form RDD for poor and rich, respectively. They evaluate the change in tax compliance for both groups between the last semester of 2022 and the first semester of 2023. The x-axis corresponds to the properties' cadastral values from 2021. Each figure indicates the estimated effect around the cut-off point, where the p-value, taken from a robust bias-corrected inference, is indicated in brackets.

Figure 3: RDD Falsification Test using a Placebo Reform Date



Notes: This figure calculates the *own-rate* effects using a placebo reform date in the middle of 2023. Panels (a) and (b) represent the reduced form RDD for poor and rich, in which we evaluate the change in tax compliance for both groups for placebo dates between the first semester of 2023 and the last semester of 2023. The x-axis corresponds to the cadastral valuation of the properties from 2021. Each figure indicates the estimated effect around the cut-off point, where the p-value, taken from a robust bias-corrected inference, is indicated in brackets.

Figure 4: Control and Treatment Messages Example - Poor Group

(a) Control Letter Message

FAIRER AND MORE EQUITABLE TSG

The TSG **increased below inflation** as in recent years and is also now **fairer and more equitable**.

Based on your tax valuation, your TSG decreased in relation to the rest.

(b) Treatment Letter Message

FAIRER AND MORE EQUITABLE TSG

The TSG **increased below inflation** as in recent years and also now it is **fairer and more equitable**.

We applied a 30% discount on the variable component of the TSG for properties with the lowest tax valuation, and an average increase of 15% for those with the highest value.

Based on your tax valuation, your TSG decreased in relation to the rest.

LOW VALUATIONS

A 30% discount was applied to the 35% of properties with the lowest tax valuation.

MEDIUM VALUATIONS

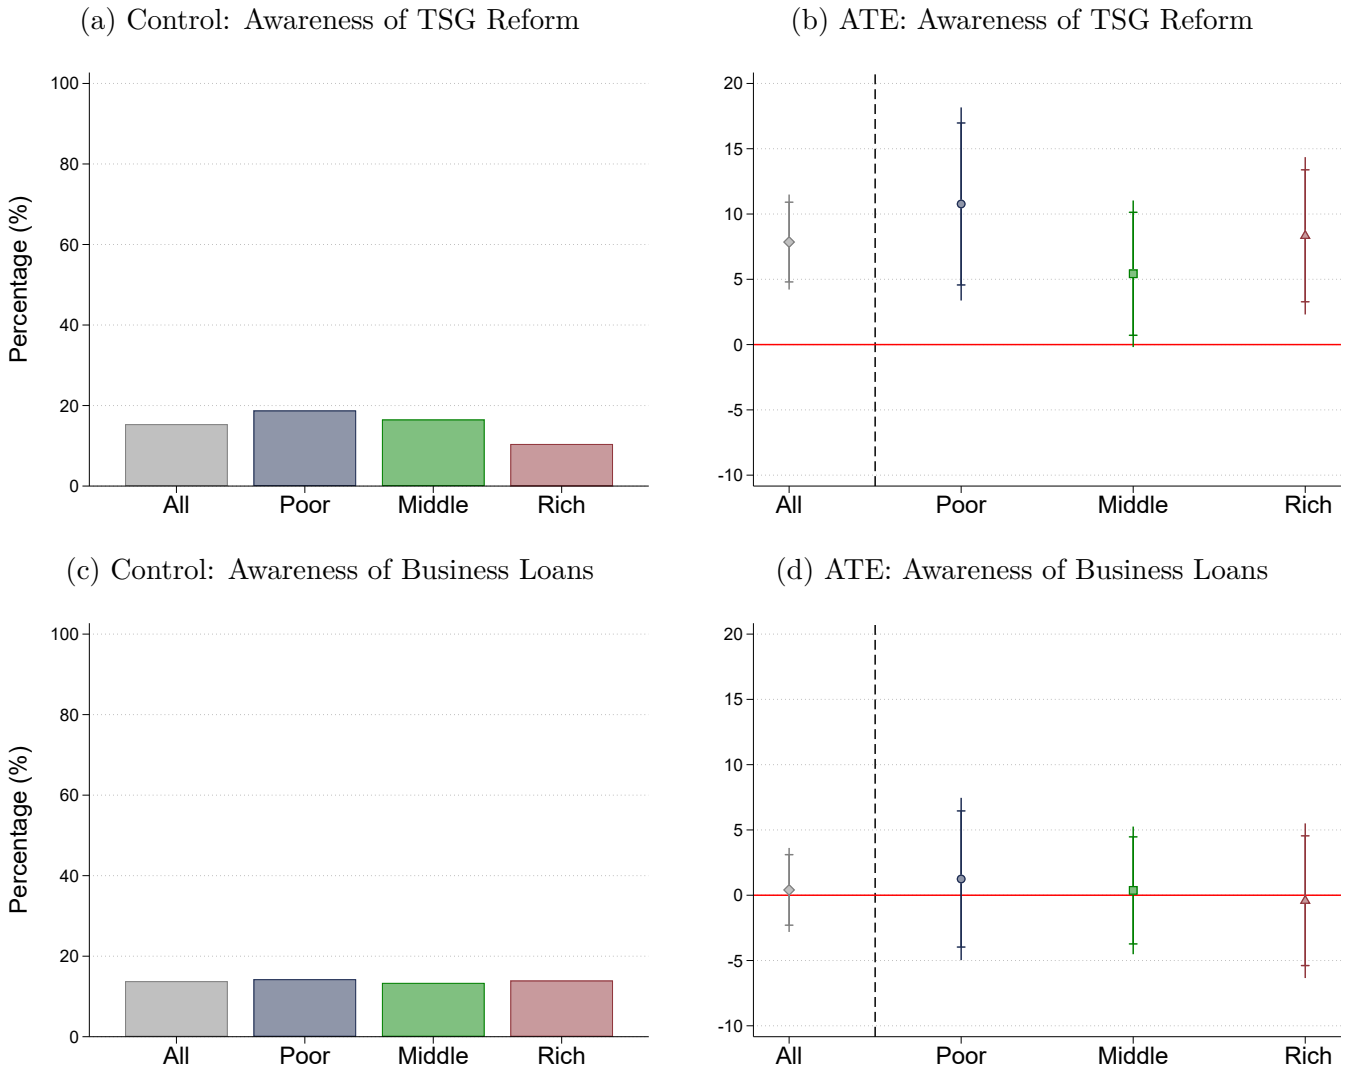
No changes were made to the next 45% of the intermediate tax valuation.

HIGH VALUATIONS

An average surcharge of 15% was made to the 20% of the highest valued properties.

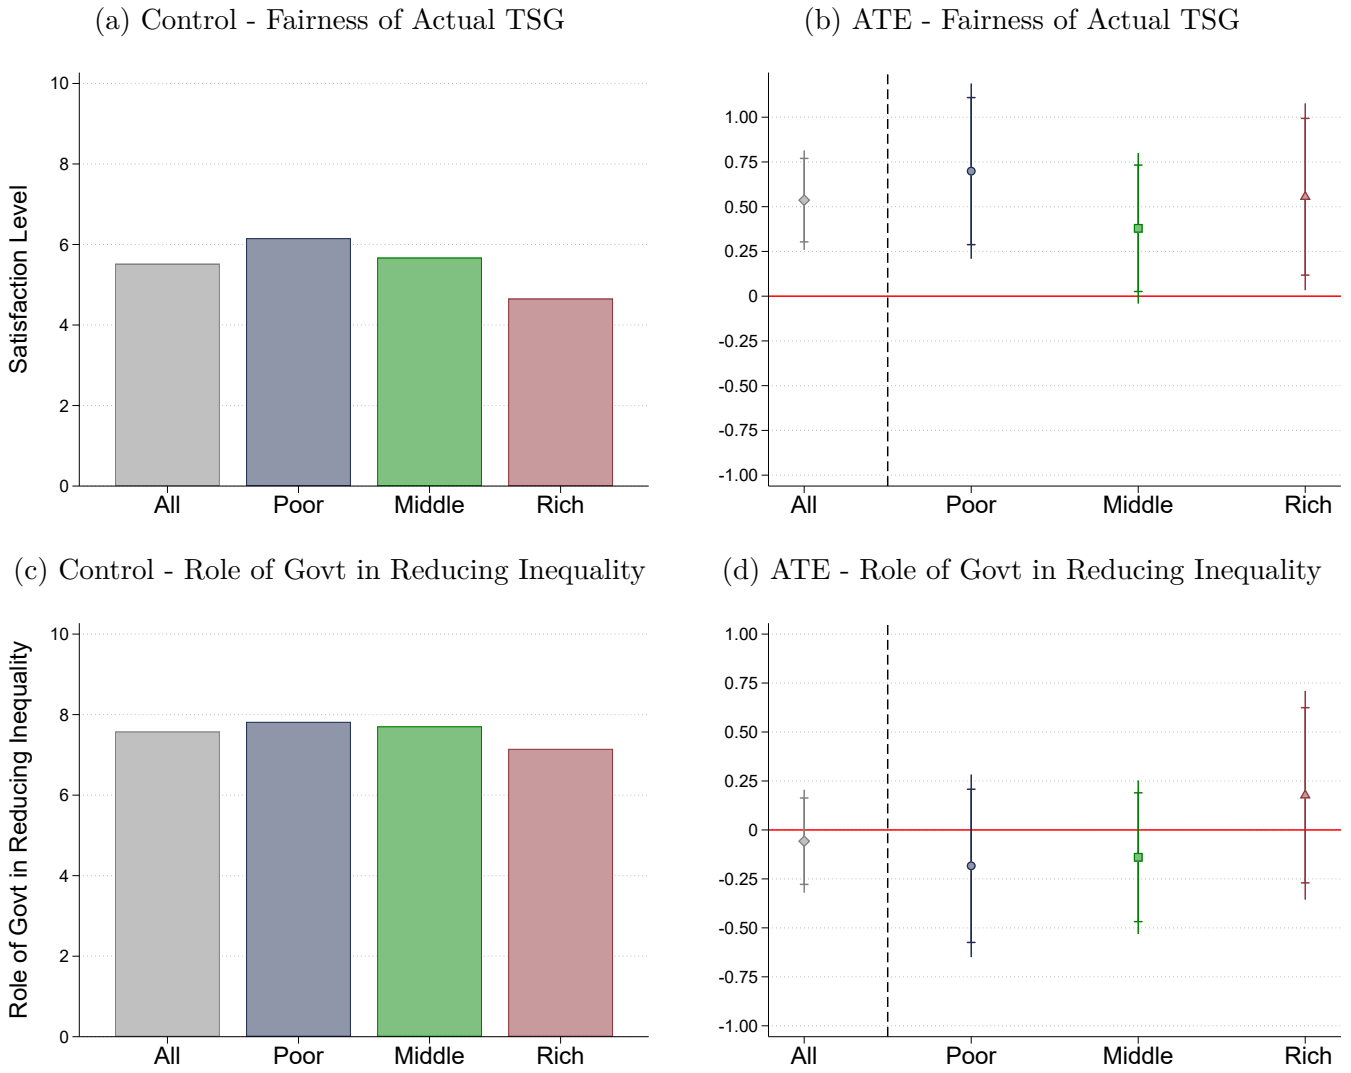
Source: Municipality of Tres de Febrero. Original letters were translated from Spanish.

Figure 5: Taxpayers' Awareness of the Reform: Survey Experiment



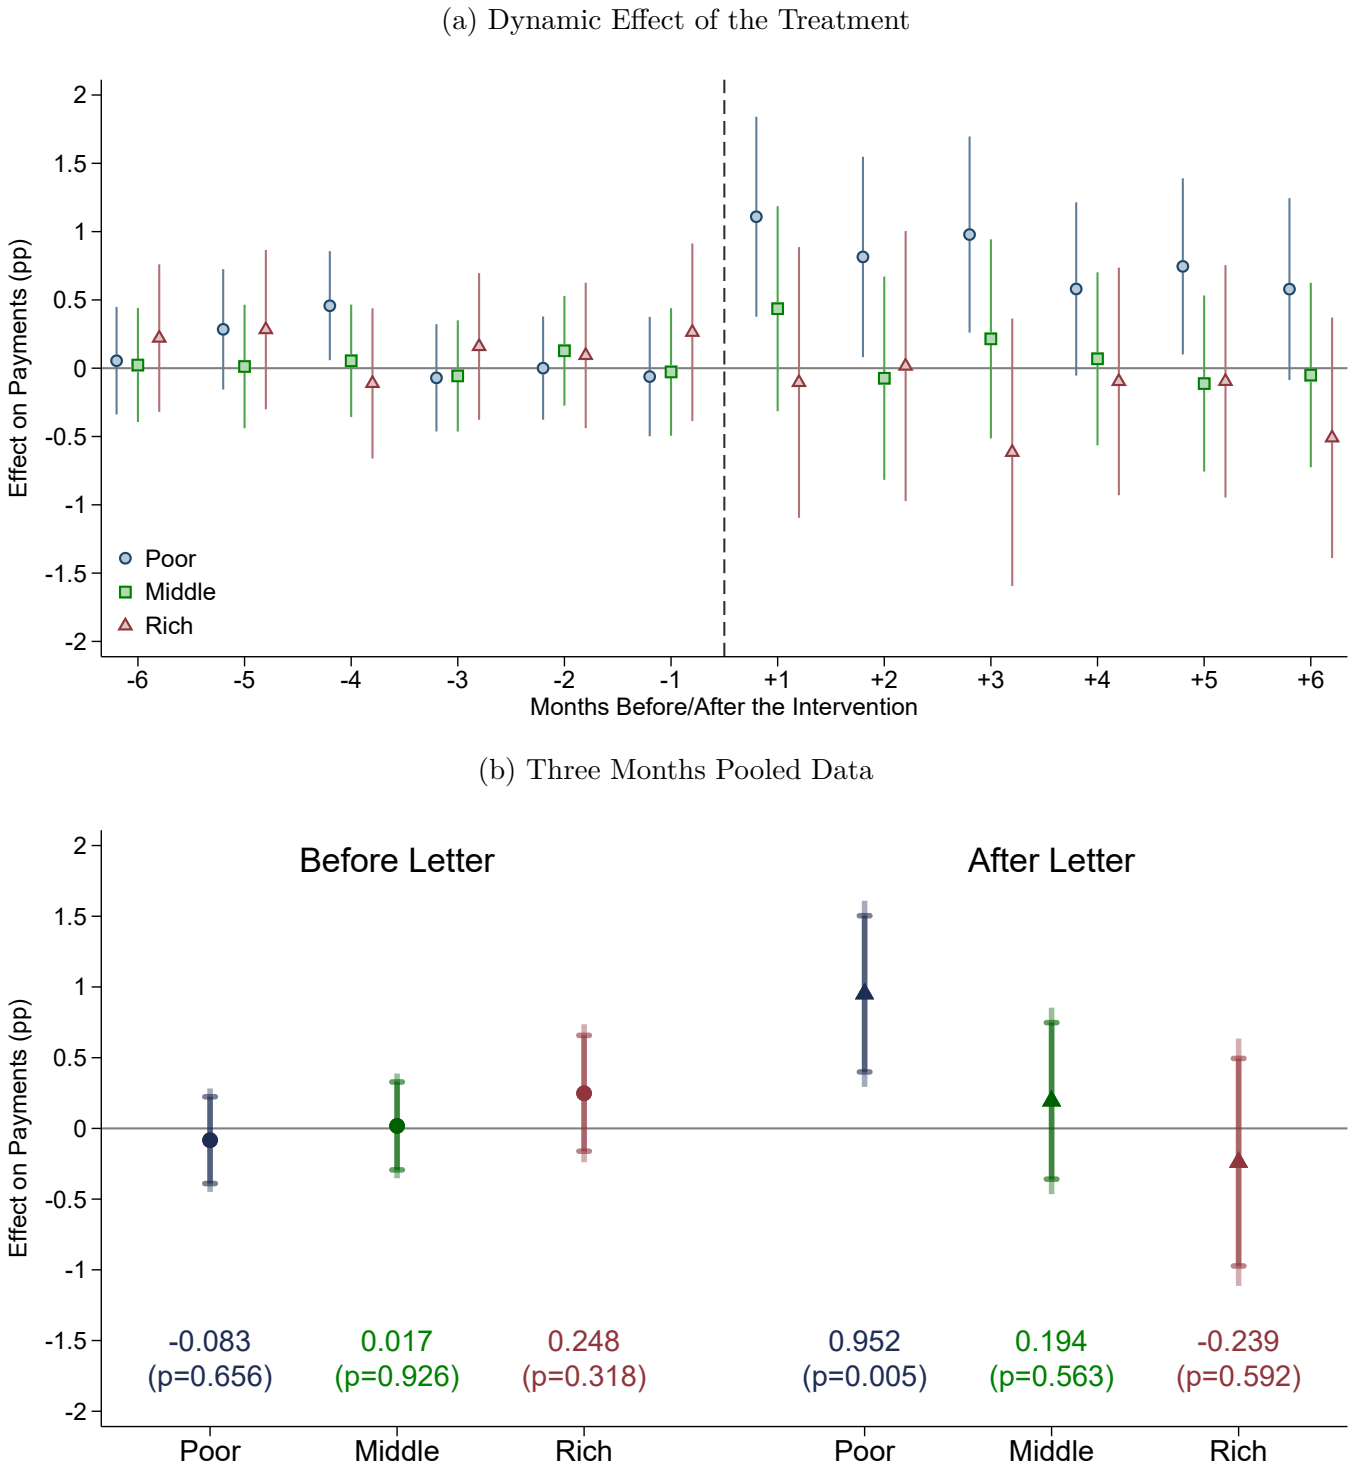
Notes: This figure uses data from our short online survey on taxpayers to assess awareness of the progressive tax reform. We invited a small subsample of subjects with email addresses on file to participate. About 2,150 households responded to the survey, for a response rate of 16.4%. The respondents received a list of recent policies and were asked to indicate which ones they had heard of. The left panels show the average knowledge of two different policies for households receiving the control letter. The right panels show the effects of the treatment letter on the awareness of the progressive reform and on the knowledge of the Tres de Febrero business loans. The vertical spikes denote 90% and 95% confidence intervals. While in the control group a minority of households are aware of the reform, the awareness is about 50% higher in treatment households.

Figure 6: The Effect on Fairness Perceptions - Survey Experiment



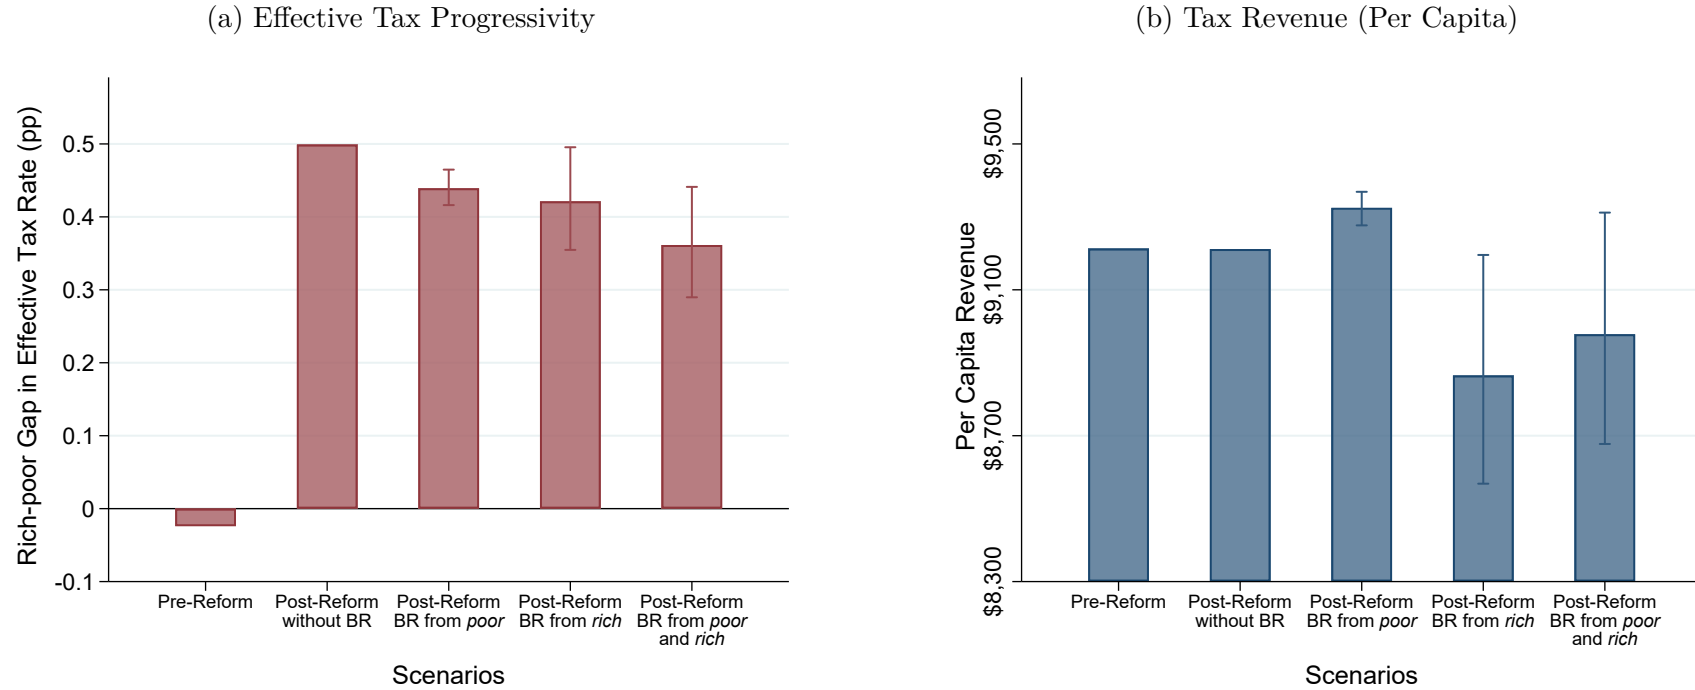
Notes: This figure uses data from our short online survey on taxpayers to estimate the impact of our progressivity information experiment on taxpayers' fairness perceptions. We asked respondents how fair they considered the distribution of the municipal tax rates between richer and poorer people, on a scale from 0 to 10 (top panels) and their perception regarding the government's role in addressing inequality (bottom panels). In both cases, a higher value implies a more progressive view of the world. The left panels show the average levels for households receiving the control letter. The right panels show the average effects of the treatment letter. The vertical spikes denote 90% and 95% confidence intervals. Panel (a) indicates relatively high perceived levels of fairness. Panel (b) shows that our information provision treatment increased this perception across the board by about 0.52 points in the 0-10 scale. Reassuringly, Panel (d) suggests that our information treatment did not affect respondents' redistribution preferences.

Figure 7: Cross-rate Effects of a Progressive Tax Reform on Own Tax Compliance



Notes: This figure compares the likelihood of paying the monthly property tax bill between treatment and control. The control group received a letter with information solely about how the reform changed the household’s own tax rate. The treatment group letter included additional information about how the reform affected the tax rates of other households (i.e., informing the progressive nature of the reform). See Figure 4 for an example. Panel (a) shows the dynamic effect of the treatment, while panel (b) shows the three months pooled effects of the treatment (pre and post treatment). ‘Poor’ (in blue) denotes households with properties valued at 750k or less, ‘Middle’ (in green) denotes properties valued between 750k and 1.5m, and ‘Rich’ (in red) denotes properties valued more than 1.5m. We estimate the coefficient for each monthly bill in separate regressions, including the 95% confidence interval for each group (and also the 90% confidence interval for the pooled estimates). The sample includes residences and shops and excludes units that made payments for 2023 before we sent the letters. We control for the past 12 months payments in each pre-treatment regression, and for the 2022 monthly payments in each post-treatment regression. Standard errors clustered at the individual taxpayer level.

Figure 8: Counterfactual Analysis: Tax Progressivity and Revenue with and without Behavioral Responses



Notes: This figure uses our key point estimates to analyze the effects of a revenue-neutral progressive reform on effective tax progressivity and tax revenues in the presence of behavioral responses (BR), and we contrast it to a counterfactual scenario where behavioral responses are muted. Panel (a) shows rich-poor effective tax rate gap, assuming different scenarios of behavioral responses (BR). Panel (b) shows the per capita revenue effects, assuming once again different scenarios of behavioral responses. The vertical spikes denote 90% confidence intervals obtained from a 5,000 repetitions bootstrap. Once behavioral responses are factored in, tax progressivity may not increase as much as intended, and a reform that was supposed to be revenue-neutral may not be so.